arXiv:1602.03885v3 [cs.GT] 12 May 2016

Accountability in Dynamic Networks

Xavier Vilaça* and Luís Rodrigues

INESC-ID, Instituto Superior Técnico, Universidade de Lisboa

A variety of distributed protocols require pairs of neighbouring nodes of a network to repeatedly interact in pairwise exchanges of messages of mutual interest. Well known examples include file-sharing systems [4] and gossip dissemination protocols [18, 17], among many other examples. These protocols can operate in very diverse settings such as wireless ad-hoc or peer-to-peer overlay networks. These settings pose three main challenges. First, networks are inherently dynamic, whether due to uncontrolled mobility or maintenance of the overlay. Second, since communication may be costly, nodes may act rationally by not sending messages, while still receiving messages from their neighbours. Third, nodes may have incomplete information about the network topology. In this work, we aim at gaining theoretical insight into how to persuade agents to exchange messages in dynamic networks.

We take a game theoretical approach to determine necessary and sufficient conditions under which we can persuade rational agents to exchange messages, by holding them accountable for deviations with punishments. Unlike previous work [19, 6, 7, 5, 20, 18, 17, 9], we do not assume the network is the result of rational behaviour or that the topology changes according to some known probability distribution, and we do not limit our analysis to one-shot pairwise interactions. We make three contributions: (1) we provide a new game theoretical model of repeated interactions in dynamic networks, where agents have incomplete information of the topology, (2) we define a new solution concept for this model, and (3) we identify necessary and sufficient conditions for enforcing accountability, i.e., for persuading agents to exchange messages in the aforementioned model.

Our results are of technical interest but also of practical relevance. We show that we cannot enforce accountability if the dynamic network does not allow for *timely punishments*. In practice, this means for instance that we cannot enforce accountability in some networks formed in file-sharing applications such as Bittorrent [4]. We also show that for applications such as secret exchange, where the benefits of the exchanges significantly surpass the communication costs, timely punishments are enough to enforce accountability. However, we cannot in general enforce accountability if agents do not possess enough information about the network topology. Nevertheless, we can enforce accountability in a wide variety of networks that satisfy 1-connectivity [15] with minimal knowledge about the network topology. This result provides a sound game theoretical foundation for the empirical results presented in [18, 17, 9], showing that we can in fact enforce accountability in gossip dissemination with connected overlay networks where nodes possess sufficient information about the network.

*E-mail: xvilaca@gsd.inesc-id.pt. Address: INESC-ID, R. Alves Redol, 9, Lisbon, Portugal. Tel.: +351 213 100 359.

1 Introduction

A variety of protocols require pairs of neighbouring nodes of a network to repeatedly interact in pairwise exchanges of messages that are of mutual interest to both parties. Well known examples include file-sharing systems [4] and gossip dissemination protocols [18, 17, 9]. These protocols can operate in very diverse settings such as wireless ad-hoc or peer-to-peer overlay networks. These settings pose three main challenges. First, networks are inherently dynamic, whether due to uncontrolled mobility or maintenance of the overlay. Second, since communication may be costly, nodes may act rationally by not sending messages, while still receiving messages from their neighbours. Third, nodes may have incomplete information about the network topology. In this work, we aim at gaining theoretical insight into how to persuade agents to exchange messages in dynamic networks.

A growing body of literature has taken a game theoretical approach [22] to address rational behaviour in a variety of distributed problems in static networks (e.g., [10, 1, 2, 3, 26]). Nodes are viewed as being under the control of rational agents [1] that seek to maximize individual utilities. In these lines, we use Game Theory to address the problem of rational behaviour in pairwise exchanges of messages over links of a dynamic network. We consider that agents obtain a benefit of receiving messages from their neighbours, and incur costs for sending and receiving messages. We are interested in protocols that satisfy two properties: (1) agents exchange messages in every pairwise interaction with neighbours, in order to provide some useful service such as file-sharing and (2) the protocols are equilibria according to some solution concept (e.g., Nash equilibrium), so that no agent gains (increases its utility) by deviating from the protocol. Protocols that satisfy these properties are said to *enforce accountability* in pairwise exchanges.

Folk Theorems. Results known as Folk Theorems have shown that it is possible to enforce accountability by holding agents accountable for deviations with punishments that decrease their utility [19], provided that (1) agents want to participate in exchanges, i.e., the benefits of receiving messages outweigh the communication costs, (2) agents perceive interactions to occur infinitely often, which is possible if the end-horizon of interactions is unknown [22], and (3) a monitoring infrastructure provides agents with sufficient information regarding the past behaviour of other agents, so that any deviation from the protocol may be promptly detected and punished. In a distributed system, agents may only learn about past behaviour of other agents in messages received from their neighbours. Therefore, a monitoring infrastructure is constrained by the network. Unfortunately, most proofs of Folk Theorems assume an exogenous monitoring infrastructure, ignoring the constraints imposed by dynamic networks. Some proofs of Folk Theorems take into account network constraints [6, 24], but they assume that networks are static or that agents know a probability distribution over topologies at each point in time. Such knowledge may not be readily available to agents in our setting. As we shall see, this poses multiple challenges not addressed by existing proofs of Folk Theorems.

Contributions. We bridge the gap between existing proofs of Folk Theorems and the goal of enforcing accountability in dynamic networks. We make three contributions: (1) we provide a new game theoretical model of repeated interactions in dynamic networks, where agents have incomplete information of the topology, (2) we define a new solution concept for this model, and (3) we identify necessary and sufficient conditions for enforcing accountability in the aforementioned model.

Game Theoretical Model. We consider that an *adversary* selects the network topologies [15]. We focus on networks that are not under the control of the agents, exemplified by wireless ad-hoc networks and overlays such as [18, 17]. To capture such exogenous restrictions on the network, we assume that the adversary is oblivious to the messages sent by agents, selecting the topologies for all times prior to the beginning of the exchanges. At each point in time, agents learn partial

[1] Henceforth, we use the designation agent when referring to both the node and the rational agents.

information about the current topology, including the identities of their current neighbours. They also possess information about the set $\mathcal{G}$ of dynamic networks that the adversary may generate. In practice, the set $\mathcal{G}$ represents basic information known by agents regarding the network structure. For instance, in an overlay network designed to disseminate data in a reliable fashion, agents expect the adversary to only generate connected graphs, whereas in a more dynamic setting such as a wireless ad-hoc network, $\mathcal{G}$ may contain a wider variety of non-connected graphs.

Solution Concept. We define a solution concept for our model, named $\mathcal{G}$-Oblivious Adversary Perfect Equilibrium ($\mathcal{G}$-OAPE). To understand its definition, it is useful to recall the notion of Nash equilibrium, which states that a protocol is an equilibrium if it is a best response, i.e., no agent can gain by deviating given that other agents also do not deviate. This definition does not suit our purposes for two reasons. First, it does not consider deviations *off the equilibrium path*, i.e., deviations that occur after histories of messages where some agents have deviated [22]. Thus, a Nash equilibrium may rely on empty threats of punishments such as punishments that also cause some utility loss to the punishers. Second, agents must know the probability distribution of the adversary generating each topology. The first issue is addressed in the literature by the notion of sequential equilibrium (SE) [14], which requires protocols to be best responses after every history of messages. To address the second issue, the notion of $\mathcal{G}$-OAPE refines SE by requiring protocols to be best responses after every history, for any fixed network in $\mathcal{G}$ generated by the adversary. In other words, we require that agents cannot gain from deviating given that the adversary may generate any network in $\mathcal{G}$.

Necessary and Sufficient Conditions. The main goal of the paper is to determine the degree of dynamism that we can tolerate in order to enforce accountability in pairwise exchanges, i.e., we aim to determine necessary and sufficient restrictions on $\mathcal{G}$ for $\mathcal{G}$-OAPE protocols that enforce accountability to exist. Our results show that the ability to enforce accountability depends not only on $\mathcal{G}$ but also on the structure and utility of pairwise exchanges.

We determined the weakest set of restrictions on $\mathcal{G}$ for which we can enforce accountability. Our first result shows that $\mathcal{G}$ must satisfy a property named *timely punishments*. Intuitively, this property states that for ever pairwise exchange between agents i and j, agent j must always be able to communicate a deviation of i to some third agent that later interacts with i, and thus is able to punish i for that deviation. This restriction is not met by some networks such as file-sharing overlays (e.g., Bittorrent [4]), where users with similar interests interact frequently with each other but only rarely with users with different interests. Our next result provides a $\mathcal{G}$-OAPE protocol that enforces accountability, assuming that $\mathcal{G}$ is restricted by timely punishments and that exchanges are *valuable*, that is, agents have a high benefit/cost ratio of receiving/sending messages and neglect download costs. Such type of exchanges occurs, for instance, when agents share small but highly valuable secrets such as private keys [18, 17].

In many cases, the assumption that pairwise exchanges are valuable is too restrictive. For instance, in file-sharing, we may expect agents to be interested in exchanging large files, but the cost of uploading such files is certainly non-negligible, and the benefit-to-cost ratio of receiving/sending files may be small. It is therefore of practical interest to understand whether we can enforce accountability in more general settings. Our final two results identify a necessary and a sufficient condition to achieve this goal. Namely, we show that $\mathcal{G}$ must satisfy, in addition to timely punishments, a property that we call *eventual distinguishability*. Roughly speaking, this property states that whenever two (or more) agents may punish i from defecting in the past (towards some other agent j), they have the information necessary to coordinate their actions, so that the total number of additional punishments for a single deviation of i is never too large. Whether $\mathcal{G}$ satisfies eventual distinguishability depends both on the properties of graphs in $\mathcal{G}$ and on the information available

(a) Valuable Pairwise Exchanges (Sec. 5) (b) General Pairwise Exchanges (Sec. 6)

Figure 1: timely punishments (TP), eventual distinguishability (ED), and connectivity (C).

to agents about the topology. Following previous work in dynamic networks [15], we consider a class of connected networks where agents have knowledge of the degree of their neighbours, which are formed by overlay networks such as [18, 17, 9]. We show that if graphs in $\mathcal{G}$ are contained in this class of connected networks, then there is a protocol that enforces accountability in general pairwise exchanges.

Summary of Results. Our results are summarized in Fig. 1. We show that, in general, timely punishments (grey area) and eventual distinguishability (dotted area) are necessary conditions for enforcing accountability, whereas connectivity is a sufficient condition (protocol $P1$). In valuable pairwise exchanges, timely punishments are sufficient for enforcing accountability (protocol $P2$).

Related work. Existing works have used Game Theory to analyse interactions in file sharing [13, 21, 16] and gossip dissemination [18, 17]. Unlike these works, we do not limit our analysis to one-shot pairwise interactions. Our proposed model and results are more related to work in models of dynamic games and network formation games. In the former, the structure of the game being repeated varies in each repetition according to a known probability distribution [19, 6, 7]. This captures repetitions of the same game where only the network topology varies. In the latter, the network is the outcome of the actions of rational agents [5, 20]. To the best of our knowledge, none of these models captures network variations as being caused by an adversary oblivious to the actions of agents, and they do not model incomplete information of the network topology. Hence, they are not appropriate for modelling unpredictable changes in the network such as physical topology changes. More importantly, they do not address the challenges of devising distributed monitoring mechanisms. Work in reputation systems such as [8] have proposed and analysed distributed systems that perform monitoring, but they do not prove Folk Theorems.

2 Model

We consider a synchronous message passing system with reliable communication in a dynamic network. The system entities are the rational agents, which send messages over links of the network, and an oblivious adversary, which selects the dynamic network at the beginning. Specifically, time is divided into rounds. Prior to the execution of the protocol, the adversary selects an evolving graph G that specifies the communication graph G^m of each round m. Let $\mathcal{G}$ be the set of all evolving graphs, and let $\mathcal{N} = \{1 \dots n\}$ be the set of agents, where n is the number of agents. As in [1, 2], we assume that edges are private: agent i can send a message to j at round m only if j is a neighbour of i in G^m, and if i sends that message, then exactly j receives it; if i has multiple neighbours, then i can discriminate neighbours by sending different messages to each. We also assume for simplicity that graphs are undirected. Finally, we assume that each agent knows n and the identities of its neighbours in each round m, prior to sending messages in m.

2.1 Pairwise Exchanges

We consider an infinitely repeated game of pairwise exchanges of messages between neighbouring agents. Given $G \in \mathcal{G}$, at each round m, every two neighbouring agents i and j have values of interest to each other. The goal is to persuade i and j to share these values plus some additional information required to monitor other agents.

Actions and histories. Agents i and j may exchange messages in one or more communication steps. For simplicity, we abstract communication by considering a finite set of individual actions, which capture most exchanges of interest, namely (1) *defection*, where i omits messages (and thus passively punishes j), (2) *cooperation*, where i sends its value plus monitoring information to j, and (3) active *punishment*, where i sends messages while causing a utility loss to j (for instance, if i sends garbage instead of the value [18]). Our results can easily be generalized to arbitrary (finite) sets of individual actions. In some of our results, we also consider that a punishment can be proportional to some constant c, i.e., can cause a utility loss proportional to c; we also consider a fourth action of *punishment avoidance*, where i avoids the cost of a punishment but does not receive the value from j. Both punishment avoidance and proportional punishments are possible actions when multiple communication steps occur between agents; later, we discuss possible implementations.

At every round m, i and j simultaneously follow an individual action, and are only informed of each other's individual actions at the end of the round. A round-m *action* a_i of i specifies the individual actions of i towards every neighbour. A round-m *action profile* $\vec{a}$ specifies the round-m actions followed by all agents. When the adversary generates an evolving graph $G \in \mathcal{G}^*$, repeated pairwise exchanges are characterized by a set $\mathcal{H}(G)$ of *histories*. A round-m history $h \in \mathcal{H}(G)$ is a pair $((\vec{a}^{m'})_{m'<m}, G)$ representing the sequence of action profiles $\vec{a}^{m'}$ followed in rounds $m' < m$ and the evolving graph G selected by the adversary. In other words, h represents global information available immediately before agents follow round-m actions. A *run* r is a function mapping each round m to a round-m history $r(m)$ [2].

Given a round-m history h and $G \in \mathcal{G}$, a protocol $\vec{\sigma}$ defines a probability distribution $Pr^{\vec{\sigma}}(r \mid G, h)$ over every run r compatible with h and G (i.e., $r(m) = h$ and G is the evolving graph in h). Let $Pr^{\vec{\sigma}}(h \mid G)$ be the probability of history h being realized conditioned on G and agents following $\vec{\sigma}$ from the beginning. We say that r is a run of $\vec{\sigma}$ after h in G if $Pr^{\vec{\sigma}}(r(m') \mid G, h) > 0$ for every round $m' > m$. As an abuse of notation, we consider that $Pr^{\vec{\sigma}}(r \mid G, h) > 0$. We say that r is a run of $\vec{\sigma}$ in G if h is the initial history.

Information and Strategies. Information available to agents can be divided into knowledge about the game structure and private observations. Regarding the former, we assume that the main information about the game structure is common knowledge [3]. We also assume that there is a subset $\mathcal{G}^* \subseteq \mathcal{G}$ of evolving graphs that the adversary may generate and that is also common knowledge. This does not preclude agents from believing that the adversary may generate any evolving graph (i.e., $\mathcal{G}^* = \mathcal{G}$); however, we show that we need minimum restrictions on $\mathcal{G}^*$ to persuade agents to exchange messages. Regarding the latter, given $G \in \mathcal{G}$ and round m, every agent i acquires some information about G^m prior to sending messages in m, which includes the identity of the round-m neighbours of i. We represent this information as a set $\mathcal{G}_i^m$ of graphs, such that $G^m \in \mathcal{G}_i^m$ and, in every graph $\bar{G}$ in $\mathcal{G}_i^m$ agent i obtains the same information about $\bar{G}^m$ and G^m. For instance, if agents only know the identity of their neighbours, then $\mathcal{G}_i^m$ is the set of graphs where i has the same set of neighbours as in G^m.

Given a round-m history h, there is a private history h_i of observations made by i up to the beginning of round m, which include the sets $\mathcal{G}_i^{m'}$ for every round $m' \leq m$, the actions of i in every round $m' < m$, and the individual actions followed by neighbours of i towards i in every round $m' < m$. We can associate to each private history h_i a round-m information set I_i containing the histories that provide the same information as in h_i. Given $G \in \mathcal{G}$, we denote by $\mathcal{I}_i(G)$ the set of information sets of i compatible with G. We make the standard assumption that agents have perfect recall, that is, agents recall all their observations made in the past. A strategy σ_i of agent i

[2] For every $m' < m$, the sequence of actions in history $r(m')$ is a prefix of the sequence of actions in $r(m)$.

[3] Every agent knows this, knows that every agent knows this, and so on.

corresponds to a distributed protocol. Specifically, for each $G \in \mathcal{G}$ and information set $I_i \in \mathcal{I}_i(G)$, $\sigma_i(\cdot \mid I_i)$ is a probability distribution over round-m actions available to i. We use the designation *protocol* for strategy profiles $\vec{\sigma}$, which specify the strategy followed by every agent.

Utility. When agents follow an action profile $\vec{a}$, agent i obtains a utility $u_i(\vec{a})$. This is the sum of the utilities obtained in interactions with each neighbour j given as the difference between the benefits of receiving a value from j and the costs of punishments and communication. Specifically, whenever j follows a cooperation or punishment action and i does not avoid a punishment, i obtains a benefit β of receiving j's value. Regarding communication costs, i incurs a normalized cost 1 of sending messages in cooperation and punishment actions, incurs no cost by defecting and avoiding punishments, and incurs the cost α of receiving messages from j. Finally, a punishment of j proportional to c causes a utility loss of $c\pi$ to i, where π is the unitary cost per punishment.

We now define the (total) expected utility $u_i(\vec{\sigma} \mid G, I_i)$ of i when agents follow a given protocol $\vec{\sigma}$, conditional on some $G \in \mathcal{G}^*$ and on a round-m information set I_i. Given a round-m history $h \in I_i$, $\vec{\sigma}$ defines a probability distribution over runs compatible with G and h. Therefore, we can compute the expected utility of i in every future round $m' \geq m$ conditional on G and h as the expected value of $u_i(\vec{a})$, where the expectation is taken relative to the probability defined by $\vec{\sigma}$ of agents following action profile $\vec{a}$ in round m'. Informally, $u_i(\vec{\sigma} \mid G, I_i)$ is the expected value of $u_i(\vec{a})$ for every future round $m' \geq m$, as computed in round m. Formally, we need to take into account both (1) the probability of each $h \in I_i$ being realized and (2) the effect of time on the value of future utilities. Regarding (1), we consider a belief system μ, specifying the probability $\mu(h \mid G, I_i)$ that h is realized. Regarding (2), we follow the standard approach [19] of assuming that the value of future utilities decays over time with a discount factor $\delta \in (0, 1)$, so the expected utility in round m' discounted to round m is $\delta^{m'-m} u_i(\vec{a})$. Let $u_i(\vec{\sigma} \mid G, h)$ be the sum of the expected utilities $u_i(\vec{a})$ discounted to m for all future rounds $m' \geq m$, conditional on G and h; then, $u_i(\vec{\sigma} \mid G, I_i)$ is the expected value of $u_i(\vec{\sigma} \mid G, h)$, where the expectation is taken relative to $\mu(h \mid G, I_i)$.

2.2 Enforcing Accountability

We say that a protocol $\vec{\sigma}$ *enforces accountability* iff (1) in $\vec{\sigma}$, agents always cooperate by exchanging their values, until some agent deviates, and (2) $\vec{\sigma}$ is an equilibrium, so that no agent gains by deviating. Regarding (2), we need to define a *solution concept*, specifying exact conditions under which $\vec{\sigma}$ is an equilibrium.

Solution Concept. We define a new solution concept for our model, which is a refinement of sequential equilibrium (SE) [14]. In its original definition, a protocol $\vec{\sigma}^*$ is said to be a SE if there is a belief system μ^* consistent with $\vec{\sigma}^*$ (see below) such that for every agent i and information set I_i, i cannot gain by deviating from $\vec{\sigma}^*$ conditioning on I_i. We refine this definition by also conditioning on every evolving graph $G \in \mathcal{G}^*$ that may be selected by the adversary. Given this, if $\vec{\sigma}^*$ is an equilibrium, then no agent i gains by deviating from σ_i^*, regardless of the evolving graph selected by the adversary and the observations made by i. This is sufficient for ensuring that agents do not gain from not punishing other agents or not forwarding monitoring information triggering additional punishments.

Formally, we say that $\vec{\sigma}^*$ is a $\mathcal{G}^*$-Oblivious Adversary Perfect Equilibrium ($\mathcal{G}^*$-OAPE) iff there is a belief system μ^* consistent with $\vec{\sigma}^*$ and $\mathcal{G}^*$ such that, for every $G \in \mathcal{G}^*$, agent i, strategy σ_i, and information set $I_i \in \mathcal{I}_i(G)$, $u_i(\vec{\sigma}^* \mid G, I_i) \geq u_i((\sigma_i, \vec{\sigma}^*_{-i}) \mid G, I_i)$, where the expectation is taken relative to μ^*, and $(\sigma_i, \vec{\sigma}^*_{-i})$ is the protocol where only i deviates from σ_i^* by following σ_i. The definition of consistent beliefs is identical in spirit to that proposed in the definition of sequential equilibrium [14]. We say that a belief system μ^* is consistent with $\vec{\sigma}^*$ and $\mathcal{G}^*$ iff there is a sequence $\vec{\sigma}^1, \vec{\sigma}^2, \ldots$ of completely mixed protocols (attribute positive probability to every possible action at every information set) converging to $\vec{\sigma}^*$ such that for every $G \in \mathcal{G}^*$, agent i, information set $I_i \in \mathcal{I}_i$,

and $h \in \mathcal{I}_i$,

$$\mu(h \mid G, I_i) = \lim_{c \to \infty} Pr^{\vec{\sigma}^c}(h \mid G) / \sum_{h' \in I_i} Pr^{\vec{\sigma}^c}(h' \mid G).$$

Intuitively, given G, if I_i is consistent with agents following $\vec{\sigma}^*$, then i believes that agents followed $\vec{\sigma}^*$, otherwise, i believes that they followed some alternative protocol.

3 Properties of $\mathcal{G}^*$-OAPE

We prove the One-Shot-Deviation principle and a property of consistent beliefs regarding a single deviation from a $\mathcal{G}^*$-OAPE.

3.1 One-Shot-Deviation Principle

We find that the One-Shot-Deviation Principle holds for $\mathcal{G}^*$-OAPE [11]. This principle states that, for a $\vec{\sigma}^*$ to be a $\mathcal{G}^*$-OAPE, it suffices that no agent i can gain by performing a *one-shot* deviation from σ_i^* at any information information set I_i, where a one-shot deviation is a strategy where i follows some action not specified by σ_i^* and follows σ_i^* afterwards. This principle will be useful later, so we provide a proof. The proof is almost identical to [11], so we only include a sketch. Let $\vec{\sigma}^*|_{I_i,a_i}$ be the protocol identical to $\vec{\sigma}^*$, except i deterministically follows a_i at I_i.

Proposition 1. ***One-Shot-Deviation Principle.*** *A protocol $\vec{\sigma}^* \in \Sigma$ is a $\mathcal{G}^*$-OAPE if and only if there exists μ^* consistent with $\vec{\sigma}^*$ and $\mathcal{G}^*$ such that, for every $G \in \mathcal{G}^*$, round m, agent i, round-m $I_i \in \mathcal{I}_i(G)$, and actions $a_i^*, a_i \in \mathcal{A}_i(G^m)$, we have $u_i(\vec{\sigma}^*|_{I_i,a_i^*} \mid G, I_i) \geq u_i(\vec{\sigma}^*|_{I_i,a_i'} \mid G, I_i) \mid G, I_i)$.*

Proof. (Sketch) We proceed as in [11], except that we fix $G \in \mathcal{G}^*$. The implication is clear: since $\vec{\sigma}^*$ is a $\mathcal{G}^*$-OAPE, agent i cannot increase its expected utility by following $\vec{\sigma}^*|_{I_i,a_i'}$ instead of $\vec{\sigma}^*$, and $u_i(\vec{\sigma}^*|_{I_i,a_i^*} \mid G, I_i) = u_i(\vec{\sigma}^* \mid G, I_i)$. As for the reverse implication, fix round-m I_i. Suppose that the right-hand side of the proposition holds and that there is σ_i such that for some $\epsilon > 0$

$$u_i((\sigma_i, \vec{\sigma}_{-i}^*) \mid G, I_i) - u_i(\vec{\sigma}^* \mid G, I_i) = 2\epsilon. \tag{1}$$

Define $m' > m$ such that $\delta^{m'-m} yn/(1-\delta) < \epsilon$. Let σ_i' be identical to σ_i at every round-m'' information set for $m'' \leq m'$, but is identical to σ_i^* at every round-m'' information set for $m'' > m'$. It can be shown using the right-hand side and backwards induction that

$$u_i(\vec{\sigma}^* \mid G, I_i) \geq u_i((\sigma_i', \vec{\sigma}_{-i}^*) \mid G, I_i) \geq u_i((\sigma_i, \vec{\sigma}_{-i}^*) \mid G, I_i) - \epsilon.$$

This contradicts (1), proving the reverse implication and concluding the proof. □

3.2 Single Deviation

We now show that the definition of consistent belief systems implies that, if an agent i has been deviating from $\vec{\sigma}^*$ up to round m and the round-m information set I_i is consistent with only i deviating from $\vec{\sigma}^*$, then i believes that only he deviated. Fix $G \in \mathcal{G}^*$, belief system μ^* consistent with $\vec{\sigma}^*$ and $\mathcal{G}^*$, agent i, and round-m information set $I_i \in \mathcal{I}_i(G)$. We say that I_i is consistent with only i deviating from $\vec{\sigma}^*$ iff there is a strategy σ_i' and history $h \in I_j$ with $Pr^{(\sigma_i', \vec{\sigma}_{-i}^*)}(h \mid G) > 0$. Let $Pr^{\vec{\sigma}}(\vec{a} \mid G, h)$ be the probability of agents following $\vec{a}$ conditioning on h and G.

Proposition 2. *If I_i is consistent with only i deviating from $\vec{\sigma}^*$, then for all strategies σ_i' consistent with I_i and G ($Pr^{(\sigma_i', \vec{\sigma}_{-i}^*)}(h' \mid G) > 0$ for some $h' \in I_i$) and all histories $h \in I_i$ we have*

$$\mu^*(h \mid G, I_i) = \frac{Pr^{(\sigma_i', \vec{\sigma}_{-i}^*)}(h \mid G)}{\sum_{h' \in I_i} Pr^{(\sigma_i', \vec{\sigma}_{-i}^*)}(h' \mid G)}.$$

Proof. Fix round-m $h \in I_i$. Let $\vec{\sigma}^1, \vec{\sigma}^2, \ldots$ be a sequence of completely mixed protocols that converge to $\vec{\sigma}^*$ such that

$$\mu^*(h \mid G, I_i) = \lim_{c \to \infty} \frac{Pr^{\vec{\sigma}^c}(h \mid G)}{\sum_{h' \in I_i} Pr^{\vec{\sigma}^c}(h' \mid G)}. \tag{2}$$

Since agents follow each action with an independent probability, $Pr^{\vec{\sigma}^c}(h \mid G)$ is the product of a set of factors, with one factor $Pr^{\sigma_j^c}(h_j \mid G)$ per agent j corresponding to the probability defined by $\vec{\sigma}^c$ of j following the sequence of actions specified in h. Specifically, this is the product of $\sigma_j^c(a_j^{m'} \mid I_j^{m'})$ for every $m' < m$, where $a_j^{m'}$ is the round-m' action of j in h, and $I_j^{m'}$ is the round-m' information set corresponding to h. The factor $Pr^{\sigma_i^c}(h_i \mid G)$ is the same for every $h \in I_i$. Let $Pr^{\vec{\sigma}^c_{-i}}(h_{-i} \mid G)$ be the product of the factors of agents different from i, and let I_i^* be the subset of histories from I_i where only i deviates from $\vec{\sigma}^*$. Since $\vec{\sigma}^c$ converges to $\vec{\sigma}^*$, if $h' \in I_i^*$, then $\lim_{c\to\infty} Pr^{\vec{\sigma}^c_{-i}}(h'_{-i} \mid G) = Pr^{\vec{\sigma}^*}(h'_{-i} \mid G)$, and if $h' \in I_i \setminus I_i^*$, $\lim_{c\to\infty} Pr^{\vec{\sigma}^c_{-i}}(h'_{-i} \mid G) = 0$. Therefore, for every $h \in \mathcal{I}_i$, we have

$$\begin{aligned}
\mu^*(h \mid G, I_i) &= \frac{\lim_{c\to\infty} Pr^{\sigma_i^c}(h_i|G) Pr^{\vec{\sigma}^c_{-i}}(h_{-i}|G)}{\lim_{c\to\infty} \sum_{h' \in I_i} Pr^{\sigma_i^c}(h'_i|G) Pr^{\vec{\sigma}^c_{-i}}(h'_{-i}|G)} &= \\
= \lim_{c\to\infty} \frac{Pr^{\sigma_i^c}(h_i|G)}{Pr^{\sigma_i^c}(h_i|G)} \frac{Pr^{\vec{\sigma}^*_{-i}}(h_{-i}|G)}{\sum_{h' \in I_i^*} Pr^{\vec{\sigma}^*_{-i}}(h'_{-i}|G)} &= \frac{Pr^{\vec{\sigma}^*_{-i}}(h_{-i}|G)}{\sum_{h' \in I_i^*} Pr^{\vec{\sigma}^*_{-i}}(h'_{-i}|G)} &= \\
= \frac{Pr^{\sigma'_i}(h_i|G)}{Pr^{\sigma'_i}(h_i|G)} \frac{Pr^{\vec{\sigma}^*_{-i}}(h_{-i}|G)}{\sum_{h' \in I_i^*} Pr^{\vec{\sigma}^*_{-i}}(h'_{-i}|G)} &= \frac{Pr^{(\sigma'_i, \vec{\sigma}^*_{-i})}(h|G)}{\sum_{h' \in I_i} Pr^{(\sigma'_i, \vec{\sigma}^*_{-i})}(h'|G)}.
\end{aligned}$$

This proves the result. □

4 Key Concepts

In the proofs of our results, we identify multiple key concepts related to the properties of the protocols and evolving graphs. We summarize them here for future reference.

Safe-bounded protocols. It turns out that it is relevant to our results to distinguish between safe and non-safe protocols, and between bounded and non-bounded protocols. Regarding the former, a protocol $\vec{\sigma}$ is *safe* if in every interaction agents either cooperate or punish each other, thus never omitting messages. Otherwise, $\vec{\sigma}$ is *non-safe*. We show that in general non-safe protocols are not $\mathcal{G}^*$-OAPE. For this reason, we restrict some of our positive and negative results to safe protocols. Regarding the latter, a protocol $\vec{\sigma}$ is *bounded* if the duration of punishments is bounded. Specifically, every protocol can be represented as a state machine [22]; in a bounded protocol, (1) the number of states is finite, (2) some of those states are cooperation states where all agents cooperate, and (3) starting from an arbitrary state, the time of convergence to a cooperation state is bounded. Bounded protocols are simpler to analyse and more useful in practice, since they ensure that memory is bounded, which is an important requirement in dynamic networks. For these reasons, we also restrict part of the analysis to bounded protocols.

Given an agent i, $G \in \mathcal{G}^*$, round-m information set $I_i \in \mathcal{I}_i(G)$, and i-edge (j, m') with $m' > m$, $\vec{\sigma}$ defines a probability $Pr^{\vec{\sigma}}(j, m' \mid G, I_i)$ of j punishing i at round m'. Let $\mathbb{E}^{\vec{\sigma},\rho}[P_i \mid G]$ be the expected number of punishments of i in the next $\rho - 1$ rounds, defined as

$$\mathbb{E}^{\vec{\sigma},\rho}[P_i \mid G, I_i] = \sum_{i\text{-edge } (j,m'): m < m' < m+\rho} Pr^{\vec{\sigma}}(j, m' \mid G, I_i).$$

Proposition 3 shows that if beliefs are consistent with $\vec{\sigma}$ and $\mathcal{G}^*$, the expected number of punishments may be conserved over time: if agent i expects to be punished an additional number c of

times after deviating in round m, and agents follow $\vec{\sigma}^*$ afterwards, then i may still expect to be punished $c-k$ times after round $m+1$, where k is the number of round-m neighbours.

Fix a safe protocol $\vec{\sigma}$, $\rho > 0$, belief system μ^* consistent with $\vec{\sigma}$ and $\mathcal{G}^*$, agent i, $G \in \mathcal{G}^*$, strategy σ_i' where i does not deviate from σ_i in rounds $m' > m$, and round-m information set $I_i \in \mathcal{I}_i(G)$ consistent with $\vec{\sigma}' = (\sigma_i', \vec{\sigma}_{-i})$ and G, such that $\mathbb{E}^{\vec{\sigma}',\rho}[P_i \mid G, I_i] \geq \mathbb{E}^{\vec{\sigma},\rho}[P_i \mid G, I_i] + c$ for some $c > 0$. Given round-m history h and action profile $\vec{a}$, let $(h, \vec{a})$ be the round-$m+1$ history obtained by appending $\vec{a}$ to the sequence of action profiles in h.

Proposition 3. *There is round-$m+1$ $I_i' \in \mathcal{I}_i(G)$ such that*

$$\mathbb{E}^{\vec{\sigma},\rho}[P_i \mid G, I_i'] \geq \mathbb{E}^{\vec{\sigma},\rho}[P_i \mid G, I_i] + c - k,$$

where k is the number of round-m neighbours of i.

Proof. By Proposition 2, we can write

$$\begin{array}{lll}
& \mathbb{E}^{\vec{\sigma}',\rho}[P_i \mid G, I_i] & = \\
= & \sum_{h \in I_i} \mu^*(h \mid G, I_i) \sum_{\vec{a}} Pr^{\vec{\sigma}'}(\vec{a} \mid G, h)(P_i(\vec{a}) + \sum_r Pr^{\vec{\sigma}}(r \mid G, (h, \vec{a}))P_i(r, m)) & \leq \\
\leq & k + \sum_{h \in I_i} \mu^*(h \mid G, I_i) \sum_{\vec{a}} Pr^{\vec{\sigma}'}(\vec{a} \mid G, h) \sum_r Pr^{\vec{\sigma}}(r \mid G, (h, \vec{a}))P_i(r, m) & = \\
= & k + \sum_{I_i' \in \mathcal{I}_i(G, I_i, m+1)} \sum_{(h,\vec{a}) \in I_i'} Pr^{\vec{\sigma}'}(I_i' \mid G, I_i)\mu^*((h, \vec{a}) \mid I_i') \sum_r Pr^{\vec{\sigma}}(r \mid G, (h, \vec{a}))P_i(r, m) & = \\
= & k + \sum_{I_i' \in \mathcal{I}_i(G, I_i, m+1)} Pr^{\vec{\sigma}'}(I_i' \mid G, I_i)\mathbb{E}^{\vec{\sigma},\rho}[P_i \mid G, I_i'].
\end{array} \tag{3}$$

where $\mathcal{I}_i(G, I_i, m+1)$ is the set of round-$m+1$ information sets compatible with I_i (for every $I_i' \in \mathcal{I}_i(G, I_i, m+1)$ and $(h, \vec{a}) \in I_i'$, $h \in I_i$), $P_i(\vec{a}) \leq k$ is the number of punishments of i in $\vec{a}$, $P_i(r, m)$ is the equivalent number of punishments that occur after round m in run r, and $Pr^{\vec{\sigma}'}(I_i' \mid G, I_i)$ is the probability of I_i' being observed given I_i. In the second step, we used the fact that summing over $h \in I_i$ and round-m' action profile $\vec{a}$ is equivalent to summing over $I_i' \in \mathcal{I}_i(G, I_i, m+1)$ and $(h, \vec{a}) \in I_i'$, and multiplying by the probability $Pr^{\vec{\sigma}'}(I_i' \mid G, I_i)\mu^*((h, \vec{a}) \mid G, I_i')$ of I_i' being reached from I_i and the history being $(h, \vec{a})$. By Proposition 2, we have

$$\begin{array}{lll}
& Pr^{\vec{\sigma}'}(I_i' \mid G, I_i)\mu^*((h, \vec{a}) \mid G, I_i') & = \\
= & \frac{\sum_{(h',\vec{a}') \in I_i'} Pr^{\vec{\sigma}'}((h', \vec{a}')|G)}{\sum_{h' \in I_i} Pr^{\vec{\sigma}'}(h'|G)} \frac{Pr^{\vec{\sigma}'}((h,\vec{a})|G)}{\sum_{(h',\vec{a}') \in I_i'} Pr^{\vec{\sigma}'}((h', \vec{a}')|G)} & = \\
= & \frac{Pr^{\vec{\sigma}'}(h|G)}{\sum_{h' \in I_i} Pr^{\vec{\sigma}'}(h'|G)} Pr^{\vec{\sigma}'}(\vec{a} \mid G, h) & = \\
= & \mu^*(h \mid G, I_i) Pr^{\vec{\sigma}'}(\vec{a} \mid G, h).
\end{array}$$

Equation (3) implies that there is $I_i' \in \mathcal{I}_i(G, I_i, m+1)$ such that

$$\mathbb{E}^{\vec{\sigma},\rho}[P_i \mid G, I_i'] \geq \mathbb{E}^{\vec{\sigma}',\rho}[P_i \mid G, I_i] - k \geq \mathbb{E}^{\vec{\sigma},\rho}[P_i \mid G, I_i] + c - k.$$

This concludes the proof. □

Punishment Opportunities. Roughly speaking, a punishment opportunity (PO) for an interaction of agent i with j is a later interaction between an agent l and i where l has been informed of a deviation of i towards j and thus has the opportunity to punish i. This requires the existence of a temporal path (a sequence of causally influenced interactions) from j to l in the evolving graph such that i cannot interfere with information forwarded from j to l. Formally, given $G \in \mathcal{G}^*$ and agent i, a round-m i-edge is a pair (j, m) where (i, j) is an edge in G^m. We say that j causally influences l in G between m and m' [15], denoted $(j, m) \rightsquigarrow^G (l, m')$, if $m < m'$ and either $j = l$ or

there is a j-edge (o, m'') in G such that $(o, m'') \rightsquigarrow^G (l, m')$. We say that j causally influences l in G without interference from i between rounds m and m', denoted as $(j, m) \rightsquigarrow_i^G (l, m')$, if the above holds for $o \neq i$. A PO of i for (j, m) in G is an i-edge (l, m') such that $(j, m) \rightsquigarrow_i^G (l, m')$.

Evasive strategies. Given a protocol $\vec{\sigma}^*$, an evasive strategy σ_i' for agent i is a strategy where i first deviates from σ_i^* by defecting some neighbours, and then hides this deviation from as many agents as possible, for as long as possible.

Indistinguishable Evolving Graphs. We say that an evolving graph G is indistinguishable from G' to agent i at round m if i acquires the same information about G and G', regardless of the protocol followed by agents. Formally, given $G \in \mathcal{G}^*$ and round $m' \geq m$, let $\mathcal{G}_j^{m'}(G)$ be the set of round-m' graphs that provide the same information to j about the round-m' topology as $G^{m'}$, and let $C_i^{m'}(G)$ be the set of agents j such that $(j, m') \rightsquigarrow^G (i, m)$; G is indistinguishable from G' to i at m iff, for every round $m' \leq m$, $C_i^{m'}(G) = C_i^{m'}(G')$ and $\mathcal{G}_j^{m'}(G) = \mathcal{G}_j^{m'}(G')$ for all $j \in C_i^{m'}(G)$.

Proposition 4 shows the intuitive result that, if an evolving graph G is indistinguishable from G' to an agent j at round m', then the probability of j observing each round-m information set under any protocol is the same, whether the evolving graph is G or G', even when conditioning on an information set of another agent i. Given a run r, let $r_j(m')$ denote the round-m' information set I_j such that $r(m') \in I_j$.

Proposition 4. *Given $G, G' \in \mathcal{G}^*$, agents i and j, round-m $I_i \in \mathcal{I}_i(G)$, and round-m' $I_j \in \mathcal{I}_j(G)$, if G and G' are indistinguishable to j at m', then for every protocol $\vec{\sigma}$ and belief system μ^* consistent with $\vec{\sigma}$ and $\mathcal{G}^*$, we have*

$$\sum_{h \in I_i} \mu^*(h \mid I_i) \sum_{r : r_j(m) = I_j} Pr^{\vec{\sigma}}(r \mid G, h) = \sum_{h \in I_i} \mu^*(h \mid I_i) \sum_{r : r_j(m) = I_j} Pr^{\vec{\sigma}}(r \mid G', h).$$

Proof. Given a protocol $\vec{\sigma}$, round $m'' \leq m'$, $G'' \in \mathcal{G}^*$, set $S_{m''}$ of agents l such that $(l, m'') \rightsquigarrow^G (j, m')$, and round-$m''$ information set $I_{S_{m''}} = \cap I_{l \in S_{m''}}$ fixing the observations of agents in $S_{m''}$, let $Pr^{\vec{\sigma}}(I_{S_1} \mid G'') = \sum_{h \in I_{S_{m''}}} Pr^{\vec{\sigma}}(h \mid G'')$. We can show that for every round m'' and round-m'' $I_{S_{m''}}$, we have

$$Pr^{\vec{\sigma}}(I_{S_{m''}} \mid G) = Pr^{\vec{\sigma}}(I_{S_{m''}} \mid G').$$

At round 1, fix action profile $\vec{a}_{S_1}$. By definition of indistinguishable graphs, we have that all agents in S_1 have the same information about the round-1 topology in G and in G' and the same set of neighbours. Therefore, we have

$$Pr^{\vec{\sigma}^*}(I_{S_1} \mid G) = Pr^{\vec{\sigma}^*}(\vec{a}_{S_1} \mid G) = Pr^{\vec{\sigma}^*}(\vec{a}_{S_1} \mid G') = Pr^{\vec{\sigma}^*}(I_{S_1} \mid G').$$

Now, fix $I_{S_{m''}}$. By the same reasons as above, agents in $S_{m''+1}$ follow every action profile $\vec{a}_{S_{m''+1}}$ with a probability that depends only on $I_{S_{m''}}$ and $\vec{\sigma}^*$. Since the probability of $I_{S_{m''}}$ is the same, whether the evolving graph is G or G', and since $I_{S_{m''}}$ and $\vec{a}_{S_{m''+1}}$ completely determine the resulting information set $I_{S_{m''+1}}$, the inductive step holds.

By consistency of beliefs and the fact that $j \in S_{m'}$, the result follows directly from the induction. □

5 Enforcing Accountability with Weakest Adversary

We identify a necessary restriction on $\mathcal{G}$ for enforcing accountability, and provide a protocol that enforces accountability assuming only this restriction and that pairwise exchanges are valuable.

5.1 Need for Timely Punishments

The weakest restriction on $\mathcal{G}^*$ is called *timely punishments*: it says that, for some bound $\rho > 0$, and for every $G \in \mathcal{G}^*$, agent i, and i-edge (j, m), there must be a PO (l, m') of i in G for (j, m) such that $m' < m + \rho$. The need for this restriction is fairly intuitive. If the adversary is not restricted by timely punishments, then there is $G \in \mathcal{G}^*$ and an agent i such that either (1) for some i-edge (j, m), there is no PO of i in G for that i-edge, or (2) there is no limit on the time it takes between an interaction of i and a corresponding PO. In case (1), i can follow an evasive strategy to ensure that no agent capable of punishing i learns about the defection, thus never being punished. In case (2), i can delay a punishment for an arbitrarily long time; the problem here is the discount factor δ: for an arbitrarily large constant d, there is an interaction of i such that, if i defects the neighbour and later follows an evasive strategy, then i is only punished after d rounds, and the utility loss of this punishment is discounted by δ^d; for a sufficiently large d, the immediate gain of defecting outweighs the loss. Unfortunately, some real networks do not always admit timely punishments. For instance, in a file sharing application, agents with similar interest may exchange files frequently, but occasionally they may interact with agents with different interests. If agents i and j have different interests and i happens to interact with j, then j may never be able to report a defection of i to agents with interests similar to i, which may be the only timely PO's of i.

Theorem 1 shows that timely punishments are necessary to enforce accountability. In the proof, an agent i follows a single evasive strategy, which we now define. In a single evasive strategy, agent i deviates from a strategy σ_i by defecting a neighbour j and then behaves as if i had not deviated, such that the probability of neighbours of i not causally influenced by j observing each information set is the same, whether i follows σ_i' or σ_i.

Formally, we define a single evasive strategy σ_i' inductively relative to a protocol $\vec{\sigma}$. Let $\vec{\sigma}' = (\sigma_i', \vec{\sigma}_{-i})$. For every round $m' < m$, σ_i' is identical to σ_i. At every round-m $I_i \in \mathcal{I}_i(G)$, with probability $\sigma_i^*(a_i^* \mid G^1)$, i follows a_i identical to a_i^* except i defects j. After observing round-$m+1$ I_i' consistent with i following a_i, i selects I_i'' identical to I_i' except I_i'' is compatible with i following a_i^* at round m. Given $m' \geq m$, round-m' $I_i \in \mathcal{I}_i(G)$, and corresponding $I_i' \in \mathcal{I}_i(G)$ selected by i, i follows a_i with probability $\sigma_i(a_i \mid I_i')$. After observing round-$m'+1$ I_i'' consistent with i following a_i', i selects I_i''' compatible with I_i'' (same observations by i prior to round m'), and compatible with the observations made by i in round-m' in I_i''. We show that, if there is no PO of i for (j, m) in G up to round m', then every neighbour of i in rounds between m and m' makes observations with the same probability, whether i follows σ_i' or σ_i. Notice that now runs of $\vec{\sigma}''$ also specify the choices of information sets made by i prior to every round.

Fix round-m $I_i \in \mathcal{I}_i(G)$. Given history $h \in I_i$, round $m' \geq m$, set $S_{m'}$ of agents not causally influenced by j between m and m' in G, and round-m' information set $I^*_{S_{m'} \cup i} = \cap_{l \in S_{m'} \cup i} I_l^*$ fixing the observations of agents in $S_{m'}$, let $Q^{\vec{\sigma}'}(I_{S_{m'} \cup \{i\}} \mid G, h)$ be the probability of agents in $S_{m'}$ observing $I^*_{S_{m'}}$ and i selecting I_i^*. Lemma 1 shows that $Q^{\vec{\sigma}'}(I^*_{S_{m'} \cup \{i\}} \mid G, h)$ is the same as the probability of agents in $S_{m'} \cup \{i\}$ observing $I_{S_{m'} \cup \{i\}}$ when they follow $\vec{\sigma}$. Since there is no PO of i for (j, m) up to round-m', every neighbour of i is in $S_{m'}$. This implies that every neighbour of i makes the same observations between m and m' with the same probability, whether i follows σ_i' or σ_i.

Lemma 1. *If there is no PO of i in G for (j, m) up to round m', then for every $h \in I_i$, round $m'' > m$ with $m'' \leq m'$, and round-m'' information set $I_{S_{m''} \cup \{i\}}$, we have*

$$Q^{\vec{\sigma}'}(I_{S_{m''} \cup \{i\}} \mid G, h) = \sum_{h' \in I_S} Pr^{\vec{\sigma}}(h' \mid G, h). \tag{4}$$

Proof. First, consider round $m'' = m+1$ and fix h. Clearly, every agent different from i follows each round-m action with the same probability, whether i follows σ_i' or σ_i. In addition, when following

σ'_i, i selects a^*_i with probability $\sigma_i(a^*_i \mid I_i)$, and sends messages according to a^*_i except to j. Since S_{m+1} includes all agents but j and i, then every $l \in S_{m+1}$ makes each observation with the same probability whether i follows σ_i or σ'_i, and i selects each action with the same probability as if i were following σ_i.

Now, suppose that the hypothesis holds for $m'' > m$ with $m'' < m'$. Fix round-m'' $I^1_{S_{m''} \cup \{i\}}$. Since i has no PO for (j, m) up to m'', for every agent $l \in S_{m''+1} \cup \{i\}$, $l \in S_{m''} \cup \{i\}$, so l follows each action a^*_l with probability $\sigma_l(a^*_i \mid G, I^1_l)$. Moreover, every round-$m''$ neighbour of l is in $S_{m''} \cup \{i\}$, hence, given $I_{S_{m''} \cup \{i\}}$, l makes observations in round m'' (if $l = i$, i selects each information set) with the same probability, whether i follows σ_i or σ'_i. The result follows directly from the hypothesis. □

We can now prove Theorem 1.

Theorem 1. *If the adversary is not restricted by timely punishments, then there is no protocol that enforces accountability.*

Proof. The proof is by contradiction. Suppose that $\vec{\sigma}^*$ enforces accountability. Suppose also that the adversary is not restricted by timely punishments. For every $\rho > 0$, we can fix $G \in \mathcal{G}^*$, agent i, and i-edge (j, m) such that, for every PO (l, m') of i for (j, m) in G, we have $m' - m \geq \rho$. This implies that, for every round $m' \geq m$ with $m' < m + \rho$, and every i-edge (l, m') with $l \neq j$, $(j, m) \rightsquigarrow^G (l, m')$ is false, since the existence of a PO (l, m') of i for (j, m) in G with $m' < m + \rho$ is equivalent to the existence of one such (l, m') with $(j, m) \rightsquigarrow^G (l, m')$. Let ρ be such that $y\delta^\rho/(1 - \delta) < 1$, where y is the (bounded) maximum difference between utilities of a single interaction. Since σ^*_i enforces accountability, i is expected to cooperate at every round-m information set $I_i \in \mathcal{I}_i(G)$ consistent with $\vec{\sigma}^*$ and G. Here, i may follow a single evasive strategy σ'_i relative to $\vec{\sigma}^*$, G, and (j, m) . Let $\vec{\sigma}' = (\sigma'_i, \vec{\sigma}^*_{-i})$. By Lemma 1, for every round $m' < m + \rho$ with $m' \geq m$, the round-m' neighbours of i observe each round-$m' + 1$ information set with the same probability, whether i follows σ'_i or σ^*_i. Since these information sets specify the individual actions followed by and taken towards i, the expected utility of i in round m' is the same, whether agents follow $\vec{\sigma}'$ or $\vec{\sigma}^*$, except i avoids at least the cost 1 of defecting j in round m. Moreover, the maximum utility difference in every round $m' \geq m + \rho$ is yn. Therefore, we have

$$u_i(\vec{\sigma}^* \mid G, I_i) - u_i(\vec{\sigma}' \mid G, I_i) < -1 + \delta^\rho yn/(1 - \delta) < 0.$$

This contradicts the assumption that $\vec{\sigma}^*$ enforces accountability, concluding the proof. □

5.2 A $\mathcal{G}^*$-OAPE for Valuable Pairwise Exchanges with Timely Punishments

We now describe a safe-bounded protocol $\vec{\sigma}^{\text{val}}$ that enforces accountability in a setting of valuable pairwise exchanges (defined below), assuming an adversary restricted by timely punishments. In valuable pairwise exchanges, we assume that agents can perform proportional punishments and punishment avoidance actions. This implies that agents can punish each other with a cost proportional to π, and they can avoid such punishments at the expense of not receiving the value from their neighbours. We also assume that the benefit-to-cost ratio of receiving/sending a value is high. Specifically, we assume that $\beta > 1 + \alpha + \rho\pi$ and $\pi > n$, where ρ is the constant in the definition of timely punishments. In Section 5.2, we discuss exchanges that can be modelled in this way, assuming that agents exchange multiple messages per round and neglect download costs ($\alpha = 0$).

We now describe $\vec{\sigma}^{\text{val}}$. In every round, agents exchange values and monitoring information that includes accusations revealing defections. These accusations are disseminated across the network, and used by agents to adjust the cost of a punishment applied to neighbours. More precisely, let ρ be as in the definition of timely punishments. At the beginning of round m, agent i keeps for every

agent l and round $m' \in \{m-\rho \dots m-1\}$ a report indicating whether l defected some neighbour in round m'. We call a report indicating a defection an accusation. In a round-m action, agents i and j exchange their reports and values, and apply punishments proportional to the number of accusations against each other. At the end of round m, i updates its reports relative to every $l \neq j$ and $m' < m$ if j does not defect i, otherwise i emits an accusation against j for round m.

Theorem 2 shows that $\vec{\sigma}^{\text{val}}$ enforces accountability in valuable pairwise exchanges. We make two minimal assumptions: (1) the adversary is restricted by timely punishments, which as we have seen is strictly necessary, and (2) agents are *sufficiently patient*, i.e., the factor δ is sufficiently close to 1, which is a standard assumption in proofs of Folk Theorems, and is necessary for future losses of punishments to always outweigh the gains of deviating in the present (recall that future losses are discounted to the present by δ). The proof shows that a defection of i is always matched by a punishment that occurs after at most ρ rounds. i gains at most n by defecting neighbours in a round, while losing at least $\delta^{\rho}\pi$. Given that $\beta > 1 + \alpha + \rho\pi$ and $\pi > n$, if δ is sufficiently close to 1, then the loss outweighs the gain. Moreover, if i avoids the cost of a punishment and of sending and receiving messages (at most $1 + \alpha + \rho\pi$), i does not receive the value, thus losing β. Since i can never influence the reports relative to itself, i never gains by deviating.

Theorem 2. *If the adversary is restricted by timely punishments and agents are sufficiently patient, then $\vec{\sigma}^{val}$ enforces accountability in valuable pairwise exchanges.*

Proof. We show that $\vec{\sigma}^{\text{val}}$ is a $\mathcal{G}^*$-OAPE. Fix $G \in \mathcal{G}^*$, agent i, round-m, round-m information set $I_i \in \mathcal{I}_i(G)$, $h \in I_i$, and actions a_i^*, a_i' such that $\sigma_i^{\text{val}}(a_i^* | I_i) > 0$. Fix $G \in \mathcal{G}^*$, agent i, round-m, round-m information set $I_i \in \mathcal{I}_i(G)$, $h \in I_i$, and actions a_i^*, a_i' such that $\sigma_i^{\text{gen}}(a_i^* | I_i) > 0$. Let $\Delta = \mathbb{E}^{\vec{\sigma}^1}(u_i \mid h, G) - \mathbb{E}^{\vec{\sigma}^2}(u_i \mid h, G)$, where $\vec{\sigma}^1 = \vec{\sigma}^{\text{gen}}|_{I_i, a_i^*}$, $\vec{\sigma}^2 = \vec{\sigma}^{\text{gen}}|_{I_i, a_i'}$, and $\vec{\sigma}^{\text{gen}}|_{I_i, a_i}$ represents the protocol that differs from $\vec{\sigma}^{\text{gen}}$ exactly in that i deterministically follows a_i at I_i. By the One-Shot-Deviation principle, it suffices to show that $\Delta \geq 0$ to prove the result. For every $m' \geq m$, let $\Delta^{m'}$ be the difference between the expected utility of i in round m', between agents following $\vec{\sigma}^1$ and $\vec{\sigma}^2$, i.e.,

$$\Delta^{m'} = \sum_r Pr^{\vec{\sigma}^1}(r \mid G, h) u_i(\vec{a}^{m'}) - \sum_r Pr^{\vec{\sigma}^1}(r \mid G, h) u_i(\vec{a}^{m'}),$$

where $\vec{a}^{m'}$ is the round-m' action profile in r. We show the following facts, where again we say that a fact holds after i follows $\vec{a}_i'$ or $\vec{a}_i^*$ if the fact holds for all runs r^1 and r^2 in G of $\vec{\sigma}^1$ and $\vec{\sigma}^2$, respectively:

(a) *Fact 1*: If $m' \geq m$ and i follows a_i', then j has an accusation against i for round m at the end of m' iff i defects a round-m neighbour l and $(l, m) \rightsquigarrow_i^G (j, m'+1)$.

Proof. At the end of round m, an agent j has an accusation against i iff i defects j. The set of neighbours that hold accusations are exactly the set of agents j such that $(l, m) \rightsquigarrow_i^G (j, m+1)$, where l is a round-m neighbour that i defects. Continuing inductively, at the end of every subsequent round m', an agent j has an accusation if j already had that accusation or received it from $l \neq i$. Either way, by the induction hypothesis, this is only the case if i defected some round-m neighbour l and $(l, m) \rightsquigarrow_i^G (j, m'+1)$. □

(b) *Fact 2*: j has an accusation against i for $m'' \neq m$ at the end of round $m' \geq m$ after i follows a_i^* iff the same is true after i follows a_i':

Proof. When $m' = m$, agents ignore the accusations relative to i and m'' sent by i in round m, so the actions a_i^* and a_i' have the same impact on the accusations held by other agents after m'.

The same is true in every subsequent round $m' > m$, except when $m' = m''$. In this case, since i follows σ_i^{val} in round m', both after following a_i^* and a_i' in round m, i defects no neighbour and thus no agent will have an accusation against i relative to m'' at the end of m'. □

(c) *Fact 3*: If $m' > m$, then $\Delta^{m'} \geq 0$.

Proof. Since i follows σ_i^{val} after m, in round m', i either cooperates with or punishes every neighbour, and these actions have a fixed cost 1, whether i follows a_i^* or a_i'. Moreover, i always receives the values from all neighbours, obtaining a benefit β and incurring a fixed cost α. So, the utility in round m' of i following a_i' may differ from that of i following a_i^* only if the neighbours of i have a different number of accusations against i. By Facts 1 and 2, the number of accusations when i follows a_i' is at least as high as when i follows a_i^*. This implies that $\Delta^{m'} \geq 0$. □

(d) *Fact 4*: If i defects a neighbour in a_i', then there exists a round $m' < m + \rho$ such that $\Delta^{m'} \geq \pi$.

Proof. Suppose that i defects a neighbour j in a_i'. There is a PO (l, m') of i for (j, m) in G with $m' - m < \rho$, such that $(j, m) \rightsquigarrow_i^G (l, m')$. By Fact 1, l has an accusation against i for round m when i follows a_i' and no accusation relative to m when i follows a_i^*. As in the proof of Fact 3, i incurs the same costs of cooperating or punishing neighbours and obtains the same benefit and incurs the same cost of receiving values, whereas i incurs the additional cost π of being punished by l when i follows a_i'. Therefore, $\Delta^{m'} \geq \pi$. □

We conclude by showing that $\Delta \geq 0$. Suppose that in a_i' agent i avoids punishments from x neighbours and defects y neighbours. By avoiding a punishment from j, i avoids at most the cost $1 + \alpha + \rho\pi$, but loses the benefit β, thus the variation in the utility of these interactions is $\beta - 1 - \alpha - \rho\pi > 0$. By defecting j, i gains at most 1. This implies that $\Delta^m \geq -y$. By Facts 3 and 4, we have

$$\Delta \geq \delta^\rho \pi - y \geq \delta^\rho \pi - n.$$

By $\pi > n$, if δ is sufficiently close to 1, $\Delta \geq 0$, as we intended to prove. □

5.3 Examples of Valuable Pairwise Exchanges

We now provide examples of interactions that meet the restrictions of valuable pairwise exchanges. First, we consider an interaction without cryptography. Then, we show how cryptography can decrease the number of restrictions on the utility.

Without cryptography. Consider that every neighbouring agents i and j can exchange messages in three phases per round and neglect download costs, which may be the case if bandwidth is asymmetric. In phase 1, i send reports indicating the number c_j of accusations against j, and similarly j may send the equivalent number c_i of accusations against i. In phase 2, i and j send c_i and c_j penance messages that cost π each, respectively. In phase 3, they exchange their values only if both agents have sent all required information in previous phases. In this setting, agents can adjust the size of penances such that $\pi > n$. Suppose that the cost of sending phase 1, 2, and 3 messages is 1. If i sends all requested information in the first two phases, i does not avoid the punishment of sending the c_i penance messages, while avoiding at most the cost 1 of defecting j by not sending the value in phase 3. As shown by Theorem 2, i does not gain from defecting. i can avoid the punishment by not sending the requested messages in the first two phases, but in this case i does not receive the value from j in phase 3, also not gaining from this deviation.

With cryptography. The previous interactions require knowledge of β in order to appropriately define the size of a penance message, and require $\beta = \Omega(n\rho)$, making it restrictive in practice if n is large. We can mitigate these problems using the technique of delaying gratification from [18]. Instead of exchanging values in phase 3 of round r, neighbours i and j exchange the values ciphered with private keys in phase 1 along with monitoring information. In phase 2, they still send the penances and in phase 3 they exchange the keys. As before, if messages are omitted in phase 1 or 2, then the agents send no further messages. Let α^κ be the cost of sending the key. This mechanism only requires $\pi > \alpha^k$ and $\beta > 1 + n\rho\pi$. Thus, it suffices that $\beta > 1 + \epsilon$, where $\epsilon = n\rho\alpha^\kappa$. If $\alpha^\kappa \ll 1$, then $1+\epsilon$ is close to the optimal restriction of $\beta > 1$. It is also possible to adjust π without knowing β: we can define the size of a penance to be larger than that of a key, but smaller than the value. The arguments that show that this mechanism is a $\mathcal{G}^*$-OAPE are the same as Theorem 2.

6 Enforcing Accountability in General Pairwise Exchanges

We now address the problem of enforcing accountability without making the assumption that pairwise exchanges are valuable. We consider the least restrictive assumptions about utility and individual actions available to agents. First, we consider the smallest benefit-to-cost ratio of receiving/sending values. We still need the benefit β to be larger than the total costs $1+\alpha$ of sending and receiving messages, or else agents would not have incentives to engage in exchanges. Second, we do not assume that agents neglect download costs ($\alpha \geq 0$), nor that they communicate in multiple steps. Therefore, proportional punishments and punishment avoidance are not available actions. We still assume that agents can defect, cooperate, or (actively) punish other agents. Since there is a trivial one-shot implementation of active punishments, where an agent sends garbage of the size of the value instead of the value, we consider that a punishment causes a utility loss of $\pi \geq \beta$.

With this in mind, we say that a protocol enforces accountability in general pairwise exchanges if it is an equilibrium for all utilities such that $\beta > 1 + \alpha$, and only has agents following the three aforementioned actions. First, we show that non-safe protocols cannot in general enforce accountability. Then, we identify a necessary and a sufficient condition for enforcing accountability with safe-bounded protocols in general pairwise exchanges. The results are generalized for non-bounded protocols in the full paper.

6.1 Need for Safe Strategies

Existing proofs of Folk Theorems devise protocols where agents provide the minimax utility in punishments [19]. In pairwise exchanges, the minimax utility is 0, and is obtained when neighbouring agents defect each other. With such strategies, deviations during punishments provide no gain and hence require no further punishments. We identify two scenarios that illustrate problems of applying minimax punishments in dynamic networks, which occur in both types of exchanges considered in this paper. The first scenario occurs when an agent i deviates and in a later interaction with j agent i does not know whether j will punish i (in which case i is allowed to defect) or expects i not to defect and triggers additional punishments if i defects. In the second scenario, an agent j knows that a defection between l and i will occur, and j knows that if he defects another agent, then any monitoring information triggering future punishments will be lost once l defects i. Although these scenarios do not imply that we can never use non-safe strategies, they suggest that similar problems may arise in general, thus non-safe strategies may require complicated mechanisms for distinguishing between exchanges where there may and may not be defections. Since it is not clear that such distinction can be made with bounded memory (agents may need to distinguish between evolving graphs), we focus on safe strategies in the remainder of the paper.

6.1.1 Ambiguous Punishment Opportunities

Consider the scenario depicted in Fig. 2. There are five agents numbered 1 to 5, and two evolving graphs G and G' from $\mathcal{G}^*$. In round 1, agent 1 interacts with agent 2, both in G and G'.

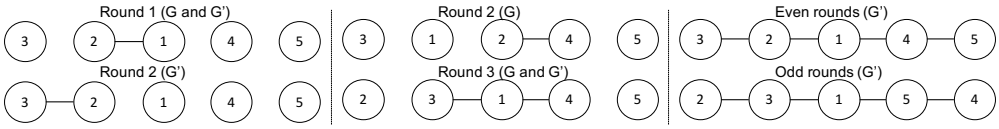

Figure 2: Ambiguous PO - G and G' are indistinguishable to 1 at round 3.

In round 2, agent 2 interacts with agent 4 in G, whereas 2 interacts with agent 3 in G'. Finally, in round 3 agent 1 interacts with agents 3 and 4 both in G and G'. Suppose that, if the adversary selects G, then 4 is the only agent capable of punishing a defection of 1 towards 2 in a timely fashion, and consider that after round 3, G' forms a line where the edges to agent 1 always form a cut between $\mathcal{N}_1 = \{2, 3\}$ and $\mathcal{N}_2 = \{4, 5\}$, and the set of neighbours of 1 alternates between $\{3, 4\}$ and $\{2, 5\}$. Notice that G' admits timely PO's for every agent i and i-edge. If the only information about the graph topology available to agents in each round is the set of their neighbours, then 1 cannot distinguish between G and G' at round 3. Suppose that there is a protocol $\vec{\sigma}^*$ that enforces accountability and requires a mutual defection between 1 and 4 if 1 defects 2. In G, since $\vec{\sigma}^*$ is a $\mathcal{G}^*$-OAPE, 1 obtains a higher utility for defecting 4. Now, consider that the adversary selects G', and consider the two cases where (a) 1 defects 2 in round 2 and defects 4 because 1 is following σ_1^*, and (b) 1 does not deviate from σ_1^* in round 2 and deviates in round 4 by defecting 4. Agent 1 can devise an evasive strategy where 1 cooperates with 4 and 5 and hides the defection towards 2 from them, such that agents 4 and 5 never punish 1. On the other hand, if i follows σ_i^* and defects 4 in round 3, then either 4 or 5 must later punish 1, because they cannot distinguish between cases (a) and (b), and in case (b) they are the only agents capable of punishing 1 for the defection. This implies that 1 obtains a strictly higher utility of cooperating with 4 rather than following σ_i^* and defecting 4, contradicting the assumption that $\vec{\sigma}^*$ enforces accountability.

Formally, given agents i, j and $G, G' \in \mathcal{G}^*$, we say that there is a round-m partition of G' regarding i and j if there is a partition $(\mathcal{N}_1, \mathcal{N}_2)$ of $\mathcal{N} \setminus \{i\}$ satisfying (1) $j \in \mathcal{N}_2$, (2) for every two i-edges (l, m') and (o, m'') in G', $(l, m') \rightsquigarrow_i^{G'} (o, m'')$ only if l and o are in the same set of the partition, and (3) for every i-edge (l, m') with $m' < m$, we have $l \in \mathcal{N}_1$. We say that i-edge (j, m) in $G \in \mathcal{G}^*$ is an ambiguous PO iff there is $G' \in \mathcal{G}^*$ such that G' is indistinguishable to i from G at m and there is a round-m partition $(\mathcal{N}_1, \mathcal{N}_2)$ of G' regarding i and j. We say that $\vec{\sigma}^*$ has ambiguous defections iff there is $G \in \mathcal{G}^*$, agent i, and ambiguous i-edge (j, m) in G such that σ_i^* requires i to defect j at some information set $I_i \in \mathcal{I}_i(G)$ consistent with only i deviating from $\vec{\sigma}^*$, i.e., there is $\sigma_i' \neq \sigma_i^*$ such that I_i is consistent with $(\sigma_i', \vec{\sigma}_{-i}^*)$ and G.

Theorem 3 proves that a $\vec{\sigma}^*$ that has ambiguous defections cannot be a $\mathcal{G}^*$-OAPE. The proof uses the notion of dual evasive strategy, which we now define. In a dual evasive strategy σ_i'', i cooperates with an agent j in round m where (j, m) is an ambiguous PO, and tries to convince agents from the two disjoint sets $\mathcal{N}_1$ and $\mathcal{N}_2$ that i is following two different strategies σ_i' and σ_i^*, respectively, where, in σ_i', i defects some agents from $\mathcal{N}_1$ and follows σ_i^* afterwards. This way, i hides the defections from the agents in $\mathcal{N}_2$, while not revealing to agents in $\mathcal{N}_1$ that i is hiding the defections from agents in $\mathcal{N}_2$. The idea of σ_i'' is similar to a single evasive strategy: after round m, i sends messages to agents in $\mathcal{N}_2$ as if i did not defect any agent, and sends messages to agents in $\mathcal{N}_1$ as if i were following σ_i''.

Specifically, suppose that i-edge (j, m) is an ambiguous PO, and fix corresponding G, G', $\vec{\sigma}^*$, $\vec{\sigma}'$, and I_i, such that $\vec{\sigma}^*$ has an ambiguous defection, I_i is a round-m information set consistent with $\vec{\sigma}'$ and G, and in $\vec{\sigma}'$ only deviates from σ_i^*, and only deviates before round m. We define a dual evasive strategy σ_i'' as follows. For every round $m' < m$ or information set I_i' incompatible with I_i, $\vec{\sigma}''$ is identical to $\vec{\sigma}^*$. The definition of σ_i'' for the remaining information sets is inductive. At every round $m' \geq m$, σ_i'' specifies a probability distribution over a pair of round-m histories (h^1, h^2) and

round-m actions followed by i. Intuitively, h^1 is used to emulate the behaviour of agents in $\mathcal{N}_1$ and h^2 is used to emulate the behaviour of agents in $\mathcal{N}_2$.

Consider round m. Since G' is indistinguishable from G at m, we have $I_i \in \mathcal{I}_i(G')$ and I_i is consistent with $\vec{\sigma}'$ and G'. Agent i selects a round-m history $h^1 \in I_i$ with probability $Pr^{\vec{\sigma}^*}(h^1 \mid G')/\sum_{h' \in I_i'} Pr^{\vec{\sigma}^*}(h' \mid G)$, where I_i' is some round-m information set consistent with $\vec{\sigma}^*$ and G. In addition, i selects h^2 with probability $Pr^{\vec{\sigma}'}(h^2 \mid G')/\sum_{h' \in I_i} Pr^{\vec{\sigma}^*}(h' \mid G')$. Now, consider round $m' \geq m$, and suppose that $\vec{\sigma}''$ is defined for every previous round. Fix a pair (h^1, h^2) selected at the end of $m'-1$. At round m', agent i does the following. First, i selects two action profiles $\vec{a}^1_{\mathcal{N}_2 \cup \{i\}}$ and $\vec{a}^2_{\mathcal{N}_1 \cup \{i\}}$ with probabilities $Pr^{\vec{\sigma}^*}(\vec{a}^1_{\mathcal{N}_2 \cup \{i\}} \mid G', h^2)$ and $Pr^{\vec{\sigma}^*}(\vec{a}^2_{\mathcal{N}_1 \cup \{i\}} \mid G', h^1)$, respectively. These are the actions that agents in $\mathcal{N}_2$ and $\mathcal{N}_1$ could have followed had i followed σ_i^* at I_i and σ_i^* from the beginning, respectively. Then, i follows individual actions towards agents in $\mathcal{N}_1$ according to a_i^1 and follows individual actions towards agents in $\mathcal{N}_2$ according to a_i^2. After observing round-$m'+1$ information set I_i', i deterministically selects two action profiles $\vec{a}^1_{\mathcal{N}_1}$ and $\vec{a}^2_{\mathcal{N}_2}$ such that the round-m' individual actions followed by agents in $\mathcal{N}_1$ towards i in I_i' are the same as those followed by these agents in $\vec{a}^1_{\mathcal{N}_1}$, and the same holds for agents in $\mathcal{N}_2$ and $\vec{a}^2_{\mathcal{N}_2}$. Let $\vec{a}^1 = (\vec{a}^1_{\mathcal{N}_1}, \vec{a}^1_{\mathcal{N}_2 \cup \{i\}}, \vec{a}^1_{\mathcal{N}_2})$ and $\vec{a}^2 = (\vec{a}^2_{\mathcal{N}_1 \cup \{i\}}, \vec{a}^2_{\mathcal{N}_2})$. The selected round-$m'+1$ histories are $(h^1, \vec{a}^2)$ and $(h^2, \vec{a}^1)$, where $(h', \vec{a}')$ represents the history that results from appending $\vec{a}'$ to the sequence of action profiles in h'. This concludes the definition.

We prove Lemma 2, which shows that by following $\vec{\sigma}''$, agent i leads agents in $\mathcal{N}_1$ to believe that i follows σ_i^* at I_i, and lead agents in $\mathcal{N}_2$ to believe that all agents followed $\vec{\sigma}^*$ since the beginning. The proof is similar to that of Lemma 1.

Given round-m' history h, let $Q_1^{\vec{\sigma}''}(h \mid G', I_i)$ be the probability of h specifying the actions followed by agents in $\mathcal{N}_2$ conditional on I_i, and the actions of other agents in h being the same as in the round-m history h^1 selected by i. Let $Q_2^{\vec{\sigma}''}(h \mid G', I_i)$ denote the corresponding probability regarding the actions of agents in $\mathcal{N}_1$ and selection of h^2. Fix μ^* consistent with $\vec{\sigma}^*$ and $\mathcal{G}^*$. Lemma 2 shows that $Q_1^{\vec{\sigma}''}(h \mid G', I_i)$ is the same as the belief held by i that the round-m' history will be h conditional on G', some information set I_i' consistent with $\vec{\sigma}^*$ and G', and agents following $\vec{\sigma}^*$. Since h determines the observations of agents in $\mathcal{N}_2$, this implies that i believes that these agents will not detect the deviations of i. Similarly, $Q_2^{\vec{\sigma}''}(h \mid G', I_i)$ is the same as the belief held by i that the round-m' history will be h conditional on G', I_i, and agents following $\vec{\sigma}^*$, hence i believes that agents in $\mathcal{N}_1$ will not detect a deviation of i from $\vec{\sigma}^*$ at and after I_i.

Lemma 2. *For every round $m' \geq m$ and round-m' $h \in \mathcal{H}$, there is a round-m information set I_i' consistent with $\vec{\sigma}^*$ and G' such that*

$$Q_2^{\vec{\sigma}''}(h \mid G', I_i) = \sum_{h' \in I_i} \mu^*(h' \mid G', I_i) Pr^{\vec{\sigma}'}(h \mid G', h'),$$

$$Q_1^{\vec{\sigma}''}(h \mid G', I_i) = \sum_{h' \in I_i'} \mu^*(h' \mid G', I_i') Pr^{\vec{\sigma}^*}(h \mid G', h').$$

Proof. We first show for $Q_2^{\vec{\sigma}''}(h \mid G', I_i)$, using induction on m'.

Consider $m' = m$. By the definition of ambiguous PO (j, m), for every $l \in \mathcal{N}_1$, $o \in \mathcal{N}_2$, and $m' < m$, it is false that $(l, m') \rightsquigarrow^{G'} (o, m)$ and $(o, m') \rightsquigarrow^{G'} (l, m)$. Since i does not interact with agents in $\mathcal{N}_2$, the actions of agents in $\mathcal{N}_2$ do not depend on the actions of agents in $\mathcal{N}_1 \cup \{i\}$, and their probability is given by $\vec{\sigma}^*$. Similarly, the actions of agents in $\mathcal{N}_1$ do not depend on the actions of agents in $\mathcal{N}_2$, and their probability is given by $\vec{\sigma}'$. This implies that the belief held by i conditional on G' and I_i that the actions of agents in $\mathcal{N}_2$ are given by h is the same as the probability of these

agents following those actions, conditional on G', I_i, and agents following $\vec{\sigma}^*$; by construction of σ_i'', this is exactly the probability of i selecting those actions. Moreover, the actions of i selected by i in h are fixed by I_i; clearly, these are the same as the actions followed by i conditional on G' and I_i. Finally, by Proposition 2, the belief that agents in $\mathcal{N}_1$ follow the actions specified in h conditional on G' and I_i is exactly the probability that those agents followed those actions, conditional on G', I_i, and agents following $\vec{\sigma}'$. Therefore, we have $\mu^*(h \mid G', I_i) = Q_2^{\vec{\sigma}''}(h \mid G', I_i)$, as we intended to prove.

Now, assume the hypothesis for round m' and consider round-$m'+1$ history $(h, \vec{a})$. We have

$$Q_2^{\vec{\sigma}''}((h, \vec{a}) \mid G', I_i) = Q_2^{\vec{\sigma}''}(h \mid G', I_i) Pr^{\vec{\sigma}^*}(\vec{a} \mid G', h),$$

which directly implies the result by the hypothesis. To see this, notice that, if h represents the actions taken by agents in $\mathcal{N}_1$ and the selection of h^2, then $h_{\mathcal{N}_1 \cup \{i\}}$ specifies the observations of agents in $\mathcal{N}_1$, so they follow $\vec{a}_{\mathcal{N}_1}$ with probability $Pr^{\vec{\sigma}^*}(\vec{a}_{\mathcal{N}_1} \mid h_{\mathcal{N}_1 \cup \{i\}})$, i.e., the product of the probabilities $\sigma_l^*(a_l \mid I_l)$ for $l \in \mathcal{N}_1$, where I_l is the round-m' information set, which, by the definition of ambiguous PO, is determined by the actions $h_{\mathcal{N}_1 \cup \{i\}}$ of agents in $\mathcal{N}_1 \cup \{i\}$. Conversely, i selects $\vec{a}_{\mathcal{N}_2 \cup \{i\}}$ with probability $Pr^{\vec{\sigma}^*}(\vec{a}_{\mathcal{N}_2 \cup \{i\}} \mid G', h_{\mathcal{N}_2 \cup \{i\}})$.

The proof for $Q_1^{\vec{\sigma}''}(h \mid G', I_i)$ is identical, except we revert the roles of $\mathcal{N}_1$ and $\mathcal{N}_2$, and i initially selects I_l' consistent with $\vec{\sigma}^*$ and G'. This concludes the proof. □

We can now prove Lemma 3.

Theorem 3. *If $\vec{\sigma}^*$ has ambiguous defections, then $\vec{\sigma}^*$ does not enforce accountability.*

Proof. Fix a belief system μ^* consistent with $\vec{\sigma}^*$ and $\mathcal{G}^*$. Fix G, i, and an ambiguous (j, m). Suppose that $\sigma_i^*(a_i^D \mid I_i) > 0$, where $I_i \in \mathcal{I}_i(G)$ is a round-m information set consistent with G and a protocol $\vec{\sigma}' = (\sigma_i', \vec{\sigma}_{-i}^*)$, $a_i^D \in \mathcal{A}_i(G^m)$ is an action where i omits a message to j, and σ_i' is some strategy identical to σ_i^* in every round $m' \geq m$. Consider G' indistinguishable from G to i at m as in the definition of ambiguous PO (j, m), and corresponding partition $(\mathcal{N}_1, \mathcal{N}_2)$. Since the i-edges form a cut between agents in $\mathcal{N}_1$ and in $\mathcal{N}_2$ in G', i can control all the information that flows between agents from the two sets. More precisely, i can follow a dual evasive strategy σ_i'' such that by Lemma 2 agents in $\mathcal{N}_1$ observe each information set with the same probability, whether i follows σ_i'' and σ_i^* at I_i, and agents in $\mathcal{N}_2$ observe information sets consistent with agents following $\vec{\sigma}^*$ from the beginning. This implies that, when i follows σ_i'' at I_i and the evolving graph is G', the expected utility of i obtained in interactions with agents in $\mathcal{N}_1$ is the same as when i follows σ_i^* (identical to σ_i' at I_i). Regarding interactions with agents $l \in \mathcal{N}_2$, both i and l cooperate if i follows σ_i''. Now, suppose that for some $h \in I_i$ with $\mu^*(h \mid G', I_i) > 0$, and run r with $Pr^{\vec{\sigma}^*}(r(m'') \mid G', h) > 0$ for every $m'' > m$, an agent $l \in \mathcal{N}_2$ does not cooperate with i in r. Then, the expected utility of i following σ_i^* obtained in interactions with agents in $\mathcal{N}_2$ is strictly lower than if i always cooperated with those agents by following σ_i'', so we have

$$u_i(\vec{\sigma}^* \mid G', I_i) < u_i((\sigma_i'', \vec{\sigma}_{-i}^*) \mid G', I_i').$$

Thus, $\vec{\sigma}^*$ does not enforce accountability.

Now, consider that when i follows σ_i^* at I_i agents in $\mathcal{N}_2$ never punish i. Fix round-m information set I_i' consistent with $\vec{\sigma}^*$ and G'. Suppose that $\vec{\sigma}^*$ enforces accountability. Then, at I_i', i must cooperate with j according to σ_i^*. Given that i gains 1 by defecting j, there must be a utility loss in future interactions if i deviates from σ_i^* by defecting. However, i may follow a dual evasive strategy σ_i''' at I_i', where i is never punished. More precisely, i can behave towards agents in $\mathcal{N}_2$ as if the information set were I_i and i followed σ_i^* such that by assumption agents in $\mathcal{N}_2$ always cooperate

with i. In σ_i''', i can also behave towards agents in $\mathcal{N}_1$ as if i followed σ_i^* at I_i', such that these agents also cooperate with i. Hence, i obtains a strictly higher utility by following σ_i''' instead of σ_i^* at I_i'. Again, $\vec{\sigma}^*$ cannot enforce accountability. This concludes the proof. □

6.1.2 Unsafe Defections

Fix a bounded protocol $\vec{\sigma}^*$ that enforces accountability, and let ρ be the maximum convergence time from an arbitrary state to a cooperation state. Suppose that, in order to avoid the problem of ambiguous defections, $\vec{\sigma}^*$ never requires an agent i to defect j if j did not deviate from $\vec{\sigma}^*$. Consider a scenario with three agents numbered 1 to 3, where, agent 1 interacts with 2 and 3 in round 1, agent 2 interacts with 3 in round $\rho - 1$, and agent 3 interacts with 1 in round ρ. Suppose that (1) if 1 defects 2 or 3 in round 1, then the only timely opportunity to punish 1 is in round ρ, and (2) if 2 defects 3 in round $\rho - 1$, then monitoring information revealing this defection to agents that may punish 2 must be forwarded by 3 to 1 (e.g., due to all temporal paths crossing this edge). Consider that $\vec{\sigma}^*$ requires 3 and 1 to mutually defect in round 3 as a punishment for a defection of 1 in round 1. Notice that $\vec{\sigma}^*$ still requires 2 to not defect 3. After the defection of 1, if agent 2 knows that 3 will defect 1, regardless of whether 2 defects 3, then 2 will always gain from defecting 3, since any monitoring information revealing this defection to other agents is lost when 3 defects 1. Instead, suppose that 3 does not defect 1 after a defection of 2. Now, we may have another problem: if 3 defects 1, its future neighbours may never learn from the defection of 2, thus believing that agents followed $\vec{\sigma}^*$ after round 1. Since the system converges to a cooperation state after ρ, future neighbours of 3 may never punish 3 for defecting 1, in which case 3 gains from defecting 1, and $\vec{\sigma}^*$ cannot enforce accountability.

Formally, given ρ, agents i, j, j', and l, rounds m and m', and $G \in \mathcal{G}^*$, a *temporal path* in G from i-edge (j, m) to j'-edge (l, m') is a sequence of c tuples (j_k, j_{k+1}, m_k) such that (1) for every $k \leq c$, $j_k \neq i$, $(j, m) \rightsquigarrow^G (j_k, m_k)$, and (j_k, j_{k+1}) is an edge in G^{m_k}, and (2) $j_{c+1} = l$ and $m_c < m'$. We say that G is unsafe if there exist rounds m, m^1, m^2, agents i, j, l, i-edges (j, m), (l, m^2), and j-edge (l, m^1), such that (1) l has no interactions after round m^2 and prior to round $m + \rho$, and (2) (l, i, m^2) is in all temporal paths from every j-edge (l, m^1) to all j-edges after round m^1 and i-edges after round m^2. A protocol $\vec{\sigma}^*$ is said to avoid ambiguous punishments iff for every set S of agents, protocol $\vec{\sigma}_S$ followed by agents in S, $G \in \mathcal{G}^*$, agents $i \notin S$ and j, information set $I_i \in \mathcal{I}_i(G)$ consistent with $(\vec{\sigma}_S, \vec{\sigma}^*_{-S})$ and G, and $h \in I_i$, $\vec{\sigma}^*$ requires j to not defect i at h. Fix an unsafe $G \in \mathcal{G}^*$ and corresponding agents i, j, l and rounds m, m^1, m^2. Given a protocol $\vec{\sigma}^*$ that avoids ambiguous punishments, let σ_i' be a strategy where i deviates from σ_i^* exactly by defecting all round-m neighbours. We say that $\vec{\sigma}^*$ has an unsafe defection iff in every run r of $(\sigma_i', \vec{\sigma}^*_{-i})$ in G, agent l defects i at round m^2.

Theorem 4 shows that no bounded protocol $\vec{\sigma}^*$ that enforces accountability can have an unsafe defection if beliefs are reasonable in the following sense. μ^* provides the simplest possible explanation to deviations, by considering the smallest set of agents that deviated and the simplest deviating strategy that explains the observations. Specifically, fix $G \in \mathcal{G}^*$, agent i, and round-m information set I_i. We define $\mu^*(\cdot \mid G, I_i)$ as follows. Agent i deterministically selects the smallest set S of agents for which there is a completely mixed protocol $\vec{\sigma}'_S$ such that I_i is consistent with $(\vec{\sigma}'_S, \vec{\sigma}^*_{-S})$ and G. Given S, i deterministically selects a protocol $\vec{\sigma}'_S$ such that, for each agent $l \in S$, l deviates in the least number of rounds, and, if l deviates in round m, then l follows the same individual action towards all round-m neighbours. It is easy to see that μ^* is consistent with $\vec{\sigma}^*$ and $\mathcal{G}^*$. We just need to establish a total order regarding the protocols $(\vec{\sigma}'_S, \vec{\sigma}^*_{-S})$, from the protocol with largest S and largest number of deviations per agent, to $\vec{\sigma}^*$, using some criteria to order protocols with the same size of S and same number of deviations per agent such as the order of the identities of agents and round numbers when agents deviate. Every such protocol is completely mixed, and the

sequence converges to $\vec{\sigma}^*$.

The proof of Lemma 4 uses the notion of lenient evasive strategy, defined as follows. In a lenient evasive strategy, agent i ignores a defection of some other agent j towards i at round m, while still ensuring that no future neighbour of i notices that i ignored the defection of j. Such strategy allows i to take advantage of an unsafe defection in an interaction with l; if the strategy σ_i^* requires i to defect l iff j does not defect i, then i can gain by deviating from σ_i^* by defecting l such that no agent will punish i in the future.

Formally, consider an unsafe evolving graph $G \in \mathcal{G}^*$, and fix corresponding agents i, j, l and rounds m, m^1, m^2. Fix also a corresponding $\vec{\sigma}^*$ that has unsafe omissions, fix $\vec{\sigma}'$ where exactly i deviates from $\vec{\sigma}^*$ by defecting all round-m neighbours, and fix round-m^2 information set $I_l \in \mathcal{I}_l(G)$ consistent with protocol $\vec{\sigma}''$ and G, where in $\vec{\sigma}''$ exactly i and j deviate from $\vec{\sigma}^*$ by defecting all round-m and round-m^1 neighbours, respectively. We define a strategy σ_l''' similar to single evasive strategies, where l pretends that j did not defect at round m. Specifically, let $S_{m'}$ be the set of agents o not causally influenced in G by round-m^1 neighbours of j between m^1 and m', and let $O_{m'}$ be the set of remaining agents o for which there is an l-edge (o', m'') with $m'' > m^2$ and $(o, m') \rightsquigarrow^G (o', m'')$. Let $L_{m'} = S_{m'} \cup O_{m'} \cup \{i\}$. Let I_l' be an information set identical to I_l in terms of the actions followed agents in $L_{m'} \setminus \{i\}$, and consistent with $\vec{\sigma}'$ and G. In σ_l', at round m^2, agent l selects I_l' and follows actions with probability $\sigma_l^*(\cdot \mid I_l')$, defecting i. For every $m' > m^2$, fix round-m' information set I_l and corresponding I_l' previously selected by l. At I_l, agent l follows actions with probability distribution $\sigma_l^*(\cdot \mid I_l')$. Given round-$m'+1$ information set I_l'', l selects I_l''' compatible with I_l' and with the round-m' observations in I_l''. We prove Lemma 3, which shows that, for all rounds $m' > m^2$, both l and neighbours of l follow actions consistent with $\vec{\sigma}'$, using similar arguments to the proof of Lemma 1.

Let $\vec{\sigma}''' = (\sigma_l''', \vec{\sigma}_{-l}^*)$. Notice that runs of $\vec{\sigma}'''$ also specify the selections of information sets by l. We still use the same notation for runs. Given a run r, an information set $I_{L_{m'}}$ compatible with r specifies the observations of agents in $L_{m'} \setminus \{l\}$ and the selections of l. We show in Lemma 3 that for all runs r of $\vec{\sigma}'''$ and $m' > m^2$, and round-m' information set $I_{L_{m'}} \in \mathcal{I}_{L_{m'}}(G)$ compatible with r, $I_{L_{m'}}$ is consistent with $\vec{\sigma}'$. By the definition of unsafe G, every round-m' neighbour of i is in $L_{m'}$, so every neighbour of i believes that agents followed $\vec{\sigma}'$ after m.

Lemma 3. *Suppose that beliefs are reasonable. For every round $m' \geq m^2$, history $h \in I_l$ with $\mu^*(h \mid G, I_l) > 0$, and run r of $\vec{\sigma}''$ after h in G, the round-m' information set $I_{L_{m'}}$ compatible with r is consistent with $\vec{\sigma}'$.*

Proof. The proof is by induction on m'. Consider $m' = m^2$. Fix $h \in I_l$ with $\mu^*(h \mid G, I_l) > 0$. Since beliefs are reasonable and h is consistent with $\vec{\sigma}''$ and G, $\mu^*(h \mid G, I_i)$ is given by conditioning on agents following $\vec{\sigma}''$. Using the same arguments as in the proof of Lemma 1, we can show that for every agent $o \in S_{m'}$, the information set I_o corresponding to h is consistent with $\vec{\sigma}'$. Moreover, since every temporal path from round-m^1 j-edges to l-edges after m^2 crosses (l, i, m^2), we have $O_{m'} = \emptyset$. Finally, l selects I_l', which is consistent with $\vec{\sigma}'$ by construction. This proves the base case.

Now, suppose that the hypothesis holds for $m' \geq m^2$ and fix round-m' history h' such that there is $h \in I_l$ with $\mu^*(h \mid G, I_l) > 0$ and $Pr^{\vec{\sigma}'}(h' \mid G, h) > 0$. Let I_l' be the round-m' information set selected by l. Fix $\vec{a}'$ with $Pr^{\vec{\sigma}'}(\vec{a}' \mid G, h') > 0$. Every agent $o \in L_{m'}$ follows a_o with probability $\sigma_o^*(a_l \mid I_o)$ if $o \neq i$, where $h \in I_o$, or o follows a_o with probability $\sigma_o^*(a_o \mid I_l')$ if $o = l$. Since I_o and I_l' are consistent with $\vec{\sigma}'$ and G, $\vec{a}'$ is also consistent with $\vec{\sigma}'$ (which is equivalent to $\vec{\sigma}^*$ after m) and G. By definition of unsafe G, for every $o \in L_{m'+1} \setminus \{i\}$, the round-$m'$ neighbours of o are in $L_{m'}$, so the round-$m'+1$ information set I_o' corresponding to $(h, \vec{a})$ is consistent with $\vec{\sigma}'$ and G. In addition, l selects I_l'' compatible with I_l' and the actions of the round-m' neighbours of l in $\vec{a}$.

Since every round-m' neighbour of l is also in $L_{m'}$, I''_l is consistent with $\vec{\sigma}'$ and G. This concludes the proof. □

We can now prove Lemma 4.

Theorem 4. *If $\vec{\sigma}^*$ is bounded and has unsafe defections and beliefs are reasonable, then $\vec{\sigma}^*$ does not enforce accountability.*

Proof. Suppose that beliefs are reasonable and that $\vec{\sigma}^*$ has unsafe omissions, and fix corresponding $G \in \mathcal{G}^*$, i, j, l, m, m^1, and m^2. Fix a reasonable belief system μ^*. We show that $\vec{\sigma}^*$ is not a $\mathcal{G}^*$-OAPE with μ^*.

Let σ'_i be a strategy that differs from σ^*_i exactly in that i defects all round-m neighbours, and let $\vec{\sigma}' = (\sigma'_i, \vec{\sigma}^*_{-i})$. Since $\vec{\sigma}^*$ avoids ambiguous punishments, i must not defect j and l at round-m information sets consistent with $\vec{\sigma}^*$ and G (and with $\vec{\sigma}'$). This implies that, in every run r of $\vec{\sigma}'$ in G, prior to round m^1, both j and l know that i deviated from σ^*_i. That is, for all $o \in \{j, l\}$, round-m^1 information set $I_o \in \mathcal{I}_o(G)$ consistent with $\vec{\sigma}'$ and G, and history $h \in I_o$, agent i deviates from σ^*_i in h at round m, by defecting o. This implies that in the definition of μ^*, the set S must include i, and i must deviate by defecting all round-m neighbour.s Since I_o is consistent with $\vec{\sigma}'$ and G, the simplest explanation for I_o is that exactly i deviated by defecting all round-m neighbours, so $\mu^*(h \mid G, I_o)$ is defined by conditioning on G and agents following $\vec{\sigma}'$.

Let σ''_j be an evasive strategy where j deviates from σ^*_j by defecting l at m^1, and later j behaves as if j did not defect l. Let $\vec{\sigma}'' = (\sigma'_i, \sigma''_j, \vec{\sigma}^*_{-i,j})$. First, suppose that in every run of $\vec{\sigma}''$ agent l still defects i at round m^2. Fix round-m^1 I_j consistent with $\vec{\sigma}'$ and G. We can show that if agents follow $\vec{\sigma}''$, then for every round $m' > m^1$, j-edge (o, m'), and round-m' information set $I_o \in \mathcal{I}_o(G)$, we have

$$\sum_{h \in I_j} \mu^*(h \mid G, I_j) \sum_{r : r_o(m') = I_o} Pr^{\vec{\sigma}''}(r \mid G, h) = \sum_{h \in I_j} \mu^*(h \mid G, I_j) \sum_{r : r_o(m') = I_o} Pr^{\vec{\sigma}'}(r \mid G, h), \qquad (5)$$

where $r_o(m')$ is the round-m' information set with $r(m') \in r_o(m')$. To see this, given round $m' > m^1$, let $S_{m'}$ be the set of agents o such that $(l, m^1) \rightsquigarrow^G (o, m')$ is false. We can show using the same arguments as Lemma 1 that (5) holds for every $m' > m$ and $o \in S_{m'}$. Since G is unsafe, every neighbour of j up to round and including m^2 is in $S_{m'}$. This proves (5) for $m' \leq m^2$.

Now, let $O_{m'}$ be the set of agents o such that, for some o' and $m'' < m'$, (o', o, m'') is in a temporal path from round-m^1 j-edges to j-edges after m^1. We can show using induction that (5) also holds for every $m' > m^2$ and $o \in O_{m'}$. Since G is unsafe, for every round $m' > m^2$, l cannot be in $O_{m'}$ and every round-m' neighbour of j is in $S_{m'} \cup O_{m'}$. This implies (5) for all j-edge (o, m') with $m' > m^1$.

We now prove the induction. First, consider $m' = m^2 + 1$. By definition of unsafe G, we have $O_{m'} = \{i\}$ and $i \in S_{m'}$. As we have seen before, in runs of $\vec{\sigma}'$ and $\vec{\sigma}''$ in G, i observes each round-m^2 information with the same probability, and, since i follows σ^*_i after round m, i also follows each round-m^2 action with the same probability. The same applies to all round-m^2 neighbours of i, except l. However, since l always defects i, it holds that i observes each round-$m^2 + 1$ information set with the same probability, whether agents follow $\vec{\sigma}'$ or $\vec{\sigma}''$. This proves (5) for $m' = m^2 + 1$. Continuing inductively, for every $o \in O_{m'+1}$, by the definition of unsafe G, both o and every round-m' neighbour of o are in $S_{m'} \cup O_{m'}$, so the result follows from the same arguments as above and the hypothesis.

Now, suppose that there is a round-m^1 information set $I_l \in \mathcal{I}_l(G)$ consistent with $\vec{\sigma}''$ such that l does not defect i. In every run of $\vec{\sigma}'$ in G agent j must not defect l at m^1. This implies that at

I_l agent l knows that both i and j deviated from σ_j^*. Again, $\mu^*(\cdot \mid G, I_l)$ is defined by conditioning on $\vec{\sigma}''$. As shown in Lemma 3, l may follow a lenient evasive strategy σ_l''' such that in every round $m' > m^2$ the neighbours of l behave as if if all agents followed $\vec{\sigma}'$ after m. Since $\vec{\sigma}^*$ and $\vec{\sigma}'$ are equivalent after m, and since the system stabilizes to a cooperation state by round $m+\rho$ and l has no interactions after m^2 and prior to $m+\rho$, in every interaction after m^2 between l and a neighbour o, both l and o cooperate. Thus, the expected utility in all future interactions is $\beta - 1 - \alpha$ when l follows σ_l'''. If l follows σ_l^* at I_l, then the expected utility in all of those interactions is at most $\beta - 1 - \alpha$. Since l gains at least 1 by defecting i at I_l, l gains by deviating from σ_l^* at I_l, hence $\vec{\sigma}^*$ cannot enforce accountability. This concludes the proof. □

6.2 Need for Eventual Distinguishability

We identify a necessary restriction on $\mathcal{G}^*$ to enforce accountability with safe-bounded protocols in general pairwise exchanges, called eventual distinguishability. Later, we generalize the results to safe (possibly non-bounded) protocols. Roughly speaking, we say that the *adversary is restricted by eventual distinguishability* if for every evolving graph G in $\mathcal{G}^*$ and agent i, eventually i stops having interactions in G where i defects some neighbour j, and then expects to be punished by two or more agents for that defection, because each of those agents cannot distinguish between two evolving graphs where it should and should not punish i. If the adversary is not restricted by eventual distinguishability and agent i keeps defecting its neighbours, then the number of punishments that i expects to receive in future rounds grows unboundedly with time. Consequently, the protocol cannot be bounded, since there is no bound on the time it takes for punishments to end after agents stop deviating.

Specifically, the formal definition of eventual distinguishability has two parts: (1) a definition of *indistinguishable rounds*, in which a defection of some agent is matched by more than one future punishment, and (2) the requirement that indistinguishable rounds must eventually stop occurring. Fig. 3 depicts an example of an indistinguishable round. There are three agents numbered 1 to 3. The adversary may generate three alternative evolving graphs G^1, G^2, and G^3. In round 1, agent 1 interacts with agent 2. In round 2, agent 2 interacts with agents 3. In round 3, the interactions depend on the evolving graph: (G^1) i interacts only with 2, (G^2) i interacts only with 3, and (G^3) i interacts with both 2 and 3. In every safe protocol that enforces accountability in general pairwise exchanges, 1 must send a message to 2 in round 1, so 1 gains by defecting 2. Suppose that round 3 is the only opportunity to punish i for defecting 2 in round 1 in a timely fashion. In an equilibrium protocol, the immediate gain 1 of defecting must be lower than the future loss β of being punished (or else 1 would gain by deviating). Since β can be arbitrarily close to 1 [4], 1 must be (deterministically) punished by 2 in G^2 and must be punished by 3 in G^3 (in which case, 2 must inform 3 of the defection of 1 in round 2). Suppose that agents only know the identities of their neighbours. Then, neither agent 2 can distinguish G^3 from G^1 nor can agent 3 distinguish G^3 from G^2 at round 3. Thus, they both have to punish 1 in G^3, and the expected number of punishments must increase by 2, even though 1 only defects one neighbour. This shows that there is a round-2 information set $I_1 \in \mathcal{I}_1(G^3)$ such that, according to the information in I_1 available to 1, the expected number of punishments of 1 conditional on G^3 and I_1 is 2. If this type of interactions keeps occurring in G^3, then, for an arbitrarily large number c, there is an information set $I_1 \in \mathcal{I}_1(G^3)$ such that, basing on the information in I_1, agent 1 expects to be punished by at least c neighbours after I_1. The fact that only agent 1 deviates plays a key role here: by the definition of equilibrium, we only need to ensure that 1 never gains by defecting a neighbour at round 1 provided that other agents do not deviate afterwards. If multiple agents deviate after round 1 (e.g., agent 2 defects 3 in round 2), then 1's defection in round 1 can be forgiven, since such behaviour

[4] Recall that we only assume that $\beta > 1 + \alpha$ and α can be small.

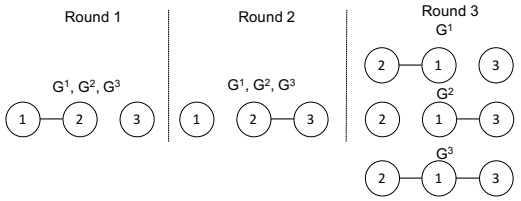

Figure 3: G^3 is indistinguishable from G^1 (G^2) to 2 (3) at round 3.

is unexpected at the time 1 decides to defect 2 in round 1. The assumption that protocols are safe is also crucial, since non-safe protocols are not susceptible to the aforementioned problem. Unfortunately, we cannot in general devise a non-safe protocol that enforces accountability.

We now provide a rigorous definition of an adversary restricted by eventual distinguishability. Given an evolving graph G, agent i, and constant ρ, let $\Phi_\rho^G(m)$ be the set of PO's (l, m') of i in G for i-edges (j, m'') such that $m'' \geq m$. We say that round m is (G, i, ρ)-indistinguishable iff there are two evolving graphs $G^1, G^2 \in \mathcal{G}^*$ such that (1) for every $G' \in \{G^1, G^2\}$ and $(j, m') \in \Phi_\rho^{G'}(m)$, G' is indistinguishable from G to j at m', (2) $|\Phi_\rho^{G^1}(m) \cap \Phi_\rho^{G^2}(m)| < k$ where k is the number of i-edges in round m, and (3) $\Phi_\rho^{G^1}(m) \cup \Phi_\rho^{G^2}(m) = \Phi_\rho^{G}(m)$. We say that the adversary is restricted by *eventual distinguishability* iff there is a constant $\rho > 0$ and a round m^* such that for all $G \in \mathcal{G}^*$ and agent i, every round $m > m^*$ is *not* (G, i, ρ)-indistinguishable.

This definition has two parts: (1) the definition of (G, i, ρ)-indistinguishable round and (2) the requirement that, eventually, no round is (G, i, ρ)-indistinguishable. We have already explained that (2) is necessary to keep punishments bounded. Part (1) is a generalization of the scenario depicted in Fig. 3 for the worst-case scenario where some agent always defects its neighbours. To understand its definition, recall that the problem identified in the example was that there were two different agents 2 and 3 who were responsible for punishing 1 in a timely fashion, after 1 defected 2 in round 1, in evolving graphs G^1 and G^2, and who could not distinguish a third evolving graph G^3 from G^1 and G^2, respectively. More generally, this happens when, for some agent i, round m, and bound ρ on the delay of punishments, (1) there are sets S_1 and S_2 of agents responsible for punishing i in G^1 and G^2, respectively, in the ρ rounds after i defects some (or all) neighbours in round m (in our example, $S_1 = \{2\}$ and $S_2 = \{3\}$), (2) S_1 and S_2 do not intersect, and (3) for $l = 1, 2$, no agent in S_l can distinguish G^3 from G^l. We further generalize this intuition by showing that the problem arises even if S_1 and S_2 intersect; they just cannot intersect in more than k agents, where k is the degree of i in round m. We also show that, if S_1 and S_2 contain all the agents that can punish i no later than ρ rounds after m, then no protocol can avoid the aforementioned problem. Given that we are considering the worst-case scenario where i always defects its neighbours, this is exactly the case when S_1 and S_2 are the agents in $\Phi_\rho^{G^1}(m)$ and $S_2 = \Phi_\rho^{G^2}(m)$, respectively.

To see this, notice that when i defects a round-m neighbour j, if (l, m') is a PO of i for (j, m), then l can clearly learn of the deviation of i and punish i in round m', so l is an agent capable of punishing i for defecting a round-m neighbour. However, since i may also defect other agents in rounds after m, these are not the only agents capable of punishing i for a defection in round m. Consider the following scenario where the neighbours of i in m and $m+1$ are j and j' respectively. Consider three PO's of i: (1) (l^1, m^1) is a PO for both (j, m) and $(j', m+1)$, (2) (l^2, m^2) is a PO for $(j', m+1)$ but not (j, m), and, similarly, (3) (l^3, m^3) is a PO for $(j', m+1)$ but not (j, m). In this case, we can devise a punishment scheme were i is punished at least once for each defection, and at most twice for defecting both j and j': (i) if i only defects j, then only l^1 punishes i; and (ii) if i defects j', then l^2 and l^3 punish i, whereas l^1 does not punish i, whether i defects j. (ii) is possible, because l^1 can be informed of the defection of i towards j'. In this case, we can say that i is punished for defecting j by l^2, even though (l^2, m^2) is not a PO of i for (j, m); it is a PO

for an i-edge from a round $m' > m$. In other words, in general, i may be punished for defecting a round-m neighbour by an agent l, in round m', if (l, m') is a PO of i for some i-edge (j', m') with $m' \geq m$. So, the set of i-edges corresponding to agents that may punish i for defecting a round-m neighbour in evolving graph G, in the ρ rounds after m, is exactly $\Phi_\rho^G(m)$.

Theorem 5 proves the need for eventual distinguishability, in order to devise safe-bounded protocols that enforce accountability in general pairwise exchanges. Fix a safe protocol $\vec{\sigma}^*$ that enforces accountability in general pairwise exchanges, fix a belief system μ^* consistent with $\vec{\sigma}^*$ and $\mathcal{G}^*$, and let ρ be the maximum time it takes for $\vec{\sigma}^*$ to converge to a cooperation state.

The proof is divided into three parts. First, we show in Lemma 4 that when agent i defects k neighbours in a round the expected number of punishments in the next $\rho - 1$ rounds must increase at least by k. Then, Lemma 6 shows that, if i constantly defects all neighbours and the adversary is not restricted by eventual distinguishability, then the expected number of punishments of i grows without bound. Finally, in the proof of Theorem 5, we identify a contradiction between Lemma 6 and the requirement of self-stabilization in bounded time.

Lemma 4. *For every $G \in \mathcal{G}^*$, agent i, information set $I_i \in \mathcal{I}_i(G)$, and protocol $\vec{\sigma}' = (\sigma_i', \vec{\sigma}^*_{-i})$ where i omits messages to k neighbours at I_i and follows σ_i^* afterwards, we have*

$$\mathbb{E}^{\vec{\sigma}',\rho}[P_i \mid G, I_i] \geq \mathbb{E}^{\vec{\sigma}^*,\rho}[P_i \mid G, I_i] + k.$$

Proof. Suppose that there is G, i, I_i, and $\vec{\sigma}'$ where i omits messages to k neighbours at I_i and

$$\mathbb{E}^{\vec{\sigma}',\rho}[P_i \mid G, I_i] < \mathbb{E}^{\vec{\sigma}^*,\rho}[P_i \mid G, I_i] + k. \tag{6}$$

At I_i, i avoids the cost k of sending messages by following σ_i' instead of σ_i^*, whereas the expected benefits and costs of receiving messages and being punished in round m are the same, whether i follows σ_i^* or σ_i'. In every round $m' > m$ and interaction with round-m' neighbour j, both i and j either cooperate with or punish each other, given that $\vec{\sigma}^*$ is safe. Thus, i incurs the fixed cost $1 + \alpha$ of sending and receiving messages, and the expected utility loss of following σ_i' instead of σ_i^* in an interaction with j is upper bounded by the probability of j punishing i times the maximum loss of β. Let S be the largest set of i-edges such that $\Delta(j, m') > 0$ for every $(j, m') \in S$, where $\Delta(j, m') = Pr^{\vec{\sigma}'}(j, m' \mid G, I_i) - Pr^{\vec{\sigma}^*}(j, m' \mid G, I_i)$. Since $\vec{\sigma}^*$ is self-stabilizing in at most ρ rounds, for every $(j, m') \in S$, we have $m' - m < \rho$, and

$$\begin{array}{lll} & u_i(\vec{\sigma}' \mid G, I_i) - u_i(\vec{\sigma}^* \mid G, I_i) & \geq \\ \geq & k - \sum_{(j',m') \in S} \delta^{m'-m} \Delta(j', m') \beta & \geq \\ \geq & k - \sum_{(j',m') \in S} \Delta(j', m') \beta & \geq \\ \geq & k - (\mathbb{E}^{\vec{\sigma}',\rho}[P_i \mid G, I_i] - \mathbb{E}^{\vec{\sigma}^*,\rho}[P_i \mid G, I_i]) \beta. & \end{array} \tag{7}$$

By (6) and the assumption that $\vec{\sigma}^*$ is independent of the utilities, there is a value of $\beta > 1$ sufficiently close to 1 such that, for some G and I_i, the utility difference between i following σ_i' and σ_i^* is strictly positive, regardless of the value of δ, so i gains by deviating from σ_i^* at I_i. This is a contradiction to $\vec{\sigma}^*$ being a $\mathcal{G}^*$-OAPE for every $\beta > 1$, thus proving the result. □

The following lemma is useful in the proof. Let σ_i' be a strategy where i defects round-m neighbours but follows σ_i^* afterwards. We show that, if (l, m') is not a PO of i in G for i-edges (j, m'') with $m'' \geq m$, then (l, m') never learns from the round-m defections of i.

Lemma 5. *For every round $m' > m$ with $m' \leq m$, set S of agents $l \neq i$ such that (l, m') is not a PO of i in G for i-edges between m and m', and round-m' information set $I_S = \cap_{l \in S} I_l$, we have*

$$\sum_{h \in I_i} \mu^*(h \mid G, I_i) \sum_{h' \in I_S} Pr^{\vec{\sigma}'}(h' \mid G, h) = \sum_{h \in I_i} \mu^*(h \mid G, I_i) \sum_{h' \in I_S} Pr^{\vec{\sigma}^*}(h' \mid G, h). \tag{8}$$

Proof. Fix m' and S. We have that for all $l \in S$ and i-edge (j,m), $(j,m) \rightsquigarrow^G (l,m')$ is false and i does not interact with l prior to m'. Given round m'', let $S_{m''}$ be the set of agents l such that (l,m'') is not a PO of i in G for i-edges between m and m''. At every round $m'' > m$ with $m'' < m'$, the information set observed by agents in $S_{m''}$, depends only on the actions of agents not in $S_{m''-1}$. Using this fact, it is direct to show the result by induction on m''. □

We can now prove Lemma 6.

Lemma 6. *If the adversary is not restricted by eventual distinguishability, then there exist $G \in \mathcal{G}^*$ and agent i such that for every $c > 0$, there is $I_i \in \mathcal{I}_i(G)$ such that*

$$\mathbb{E}^{\vec{\sigma}^*,\rho}[P_i \mid G, I_i] \geq c.$$

Proof. Suppose that the adversary is not restricted by eventual distinguishability. There is $G \in \mathcal{G}^*$, i, and an infinite sequence of rounds $m_1, m_2, \ldots$ such, that for every k, m_k is not (G,i,ρ)-distinguishable. Given round m, let $\vec{\sigma}' = (\sigma_i', \vec{\sigma}^*_{-i})$ be the protocol where i always defects up to and including round m, and follows σ_i^* afterwards. By Lemma 4, for every $G' \in \mathcal{G}^*$ and round-m information set $I_i \in \mathcal{I}_i(G)$,

$$\mathbb{E}^{\vec{\sigma}',\rho}[P_i \mid G', I_i] \geq \mathbb{E}^{\vec{\sigma}^*,\rho}[P_i \mid G', I_i] + k, \tag{9}$$

where k is the number of round-m neighbours of i. By Proposition 5, only agents agents in $\Phi_\rho^{G'}(m)$ vary the probability of punishing i, hence we have

$$\mathbb{E}^{\vec{\sigma}',\rho}[P_i \mid G', I_i] = \mathbb{E}^{\vec{\sigma}^*,\rho}[P_i \mid G', I_i] + \sum_{(j,m') \in \Phi_\rho^{G'}(m)} \Delta(j,m').$$

In particular, this is true for $m = m_k$ and $G' \in \{G^1, G^2\}$ such that (1) for every $(j,m) \in \Phi_\rho^{G'}(m)$, G' is indistinguishable from G to j at m, (2) $|\Phi_\rho^{G^1}(m) \cap \Phi_\rho^{G^2}(m)| < k$, and (3) $\Phi_\rho^{G^1}(m) \cup \Phi_\rho^{G^2}(m) = \Phi_\rho^{G}(m)$. By Proposition 4, every such j must punish i with the same probability both in G and G', because G' is indistinguishable from G to j at m, so by (9) we can write

$$\begin{array}{lll}
 & \mathbb{E}^{\vec{\sigma}',\rho}[P_i \mid G, I_i] & = \\
= & \mathbb{E}^{\vec{\sigma}^*,\rho}[P_i \mid G, I_i] + & \\
 & + \sum_{(j,m') \in \Phi_\rho^{G^1}(m)} \Delta(j,m) + \sum_{(j,m') \in \Phi_\rho^{G^2}(m)} \Delta(j,m) - & \\
 & - \sum_{(j,m') \in \Phi_\rho^{G^1}(m) \cap \Phi_\rho^{G^2}(m)} \Delta(j,m) & \geq \\
\geq & \mathbb{E}^{\vec{\sigma}^*,\rho}[P_i \mid G, I_i] + 2k - (k-1) & \geq \\
\geq & \mathbb{E}^{\vec{\sigma}^*,\rho}[P_i \mid G, I_i] + k + 1. &
\end{array}$$

By Proposition 3, there is round-$m+1$ $I_i' \in \mathcal{I}_i(G)$ such that

$$\mathbb{E}^{\vec{\sigma}',\rho}[P_i \mid G, I_i'] \geq \mathbb{E}^{\vec{\sigma}^*,\rho}[P_i \mid G, I_i] + k + 1 - k \geq \mathbb{E}^{\vec{\sigma}^*,\rho}[P_i \mid G, I_i] + 1.$$

This implies that, after every round m_k, there is an information set such that the expected number of punishments of i increases by at least 1 if i follows σ_i', while it does not decrease after every other round. It is easy to show using induction that, for every round m, there is a round-m information set $I_i \in I_i(G)$ such that the expected number of punishments of i at I_i is at least k^*, where k^* is the largest k with $m_k < m$. We can have k^* to be arbitrarily large, by arbitrarily increasing m, hence the result follows. □

Finally, we prove Theorem 5.

Theorem 5. *If the adversary is not restricted by eventual distinguishability, then there is no safe-bounded protocol that enforces accountability in general pairwise exchanges.*

Proof. The proof is by contradiction. Suppose that the adversary is not restricted by eventual distinguishability. By Lemma 6, there is G and i such that for every $c > 0$, there is a round m and round-m $I_i \in \mathcal{I}_i(G)$ such that $\mathbb{E}^{\vec{\sigma}^*,\rho}[P_i \mid G, I_i] \geq c$. This is true for m large enough such that c is larger than the number of i-edges between m and $m + \rho$. This is a contradiction to $\vec{\sigma}^*$ converging to a cooperation state in ρ rounds. □

6.3 Need for Distinguishability in Non-bounded Protocols

We now prove a slightly weaker result than the previous section, for the more general case of safe but possibly non-bounded protocols. We show that, in order to enforce accountability in general pairwise exchange protocols, the adversary must be restricted by *frequent distinguishability*, which we now define. As in the definition of eventual distinguishability, frequent distinguishability has two parts: (1) the definition of indistinguishable round, and (2) the requirement that indistinguishable rounds do not occur frequently, where by frequently we mean they keep occurring at every interval of c rounds for some constant c. Part (1) is the same as in eventual distinguishability. Part (2) differs exactly in that indistinguishable rounds may occur infinitely often, but interval between two consecutive indistinguishable rounds cannot be bounded. Formally, we say that the *adversary is restricted by frequent distinguishability* iff for every constant $c > 0$, there is $\rho > 0$ such that, for all $G \in \mathcal{G}^*$, and agent i, there is a sequence of c consecutive rounds in G that are not (G, i, ρ)-indistinguishable.

Theorem 6 proves the need for frequent distinguishability. The proof is almost identical to that of Theorem 5, in that we show by contradiction that, if the adversary is not restricted by frequent distinguishability, then the expected number of punishments of some agent is unbounded. The only difference is that, in every indistinguishable round, the expected number of punishments grows by $1 - \epsilon$ for an arbitrarily small ϵ.

We first show that, for an arbitrarily small $\epsilon > 0$, there is a sufficiently large delay ρ such that, if i deviates from a safe equilibria protocol $\vec{\sigma}^*$ at round m by first defecting k neighbours, the expected number of punishments in the next ρ rounds must be higher by at least $k - \epsilon$ than following $\vec{\sigma}^*$.

Lemma 7. *For every $\epsilon > 0$, there is $\rho > 0$ such that, for all $G \in \mathcal{G}^*$, agent i, information set $I_i \in \mathcal{I}_i(G)$, and protocol $\vec{\sigma}' = (\sigma_i', \vec{\sigma}^*_{-i})$ where i omits messages to k neighbours at I_i and follows σ_i^* afterwards, we have*

$$\mathbb{E}^{\vec{\sigma}',\rho}[P_i \mid G, I_i] \geq \mathbb{E}^{\vec{\sigma}^*,\rho}[P_i \mid G, I_i] + k - \epsilon.$$

Proof. By contradiction, if the Lemma is false, then there exists $\epsilon > 0$ such that, for an arbitrarily large ρ, we have

$$\mathbb{E}^{\vec{\sigma}',\rho}[P_i \mid G, I_i] < \mathbb{E}^{\vec{\sigma}^*,\rho}[P_i \mid G, I_i] + k - \epsilon.$$

By following σ_i', i gains k while losing at most $\delta(k - \epsilon)\beta$ in the next ρ rounds. In the rounds after $m + \rho$, i loses at most β per round. Therefore, the utility difference is at most

$$k - \delta\beta(k - \epsilon + \frac{\delta^\rho}{1 - \delta}).$$

The constant ρ can be arbitrarily large so that $\epsilon > \delta^\rho/(1 - \delta)$, since ϵ and δ are constant on ρ. If this is the case, then i gains by following σ_i', and $\vec{\sigma}^*$ cannot be a $\mathcal{G}^*$-OAPE. This is a contradiction, concluding the proof. □

We now prove that, if the adversary is not restricted by frequent distinguishability, then the expected number of punishments of some agent is unbounded.

Lemma 8. *If the adversary is not restricted by frequent distinguishability, then, for every $\rho > 0$, there is $G \in \mathcal{G}^*$ and agent i such that for every $b > 0$, there is $I_i \in \mathcal{I}_i(G)$ such that*

$$\mathbb{E}^{\vec{\sigma}^*,\rho}[P_i \mid G, I_i] \geq b.$$

Proof. Suppose that the adversary is not restricted by frequent distinguishability. There exists $c > 0$ such that, for every ρ, there is i and G such that frequently, i.e. at most every c rounds, some round is (G, i, ρ)-indistinguishable. Similar to the proof of Lemma 4, we can show using Lemma 7 that, if i keeps defecting its neighbours, then the expected number of punishments grows $1 - 2\epsilon$ every (G, i, ρ)-indistinguishable round, while decreasing at most ϵ in other rounds, for an arbitrarily small ϵ. This means that, every period of c rounds, the expected number of punishments grows

$$\epsilon' = 1 - 2\epsilon - (c-1)\epsilon = 1 - (c+1)\epsilon.$$

Since c is constant on ρ, we can find ρ large enough and ϵ small enough such that $\epsilon < 1/(c+1)$, in which case $\epsilon' > 0$. This shows that the expected number of punishments is unbounded, as we intended to prove. □

We can now conclude with the proof of Theorem 6.

Theorem 6. *If the adversary is not restricted by frequent distinguishability, then there is no safe protocol that enforces accountability in general pairwise exchanges.*

Proof. By Lemma 8, if the adversary is not restricted by frequent distinguishability, then for an arbitrarily large ρ, there is G and i such that the expected number of punishments of i in the ρ rounds following any round is unbounded. In particular, for an arbitrarily large ρ, there is $I_i \in \mathcal{I}_i(G)$ such that the expected number of punishments in the ρ rounds after I_i is larger than the number of i-edges, and i interacts with some neighbour j at I_i, so i cannot be punished for defecting j at I_i in the next ρ rounds. This is a contradiction to Lemma 7, concluding the proof. □

6.4 A $\mathcal{G}^*$-OAPE for General Pairwise Exchanges with Connectivity

We introduce a restriction of connectivity on $\mathcal{G}^*$ that ensures that the adversary is restricted by eventual distinguishability and is sufficient for devising safe protocols that enforce accountability in general pairwise exchanges. To understand this condition, it is useful to first recall the scenario of Figure 3. In this scenario, the problem arises because neither agent 2 can distinguish G^1 from G^3 at round 3 nor agent 3 can distinguish G^2 from G^3 at round 3, and thus they cannot coordinate their actions to punish 1 only once. This problem can be avoided if agents can learn the degree of their neighbours prior to deciding whether to punish them. However, the knowledge of the degree is not sufficient to satisfy eventual distinguishability, since for instance we may have a scenario similar to the one depicted in Figure 3 where agents 2 and 3 do not interact with 1 in the same round and cannot communicate with each other to coordinate the punishments. This can be avoided if it is always the case that either 2 and 3 interact in the same round with 1 or the first agent to interact with 1 can causally influence the other without interference from 1. More generally, it suffices that, for every $G \in \mathcal{G}^*$, agent i, and round m, there is ρ such that the round-m neighbours of i causally influence every round-$m+\rho$ neighbour of i between m and $m+\rho$ without interference from i. This condition is exactly met with $\rho = n$ when $\mathcal{G}^*$ is restricted by a condition similar to 1-connectivity from [15]: we say that the adversary is *restricted by connectivity* iff agents know the degree of their neighbours and, for every $G \in \mathcal{G}^*$, agent i, and round m, the graph obtained from G^m by removing

the edges to i is connected. This condition is also met by overlays for gossip dissemination such as [18, 17, 9].

We now define a safe-bounded protocol $\vec{\sigma}^{\text{gen}}$ that enforces accountability in general pairwise exchanges, assuming that the adversary is restricted by connectivity. Fix $G \in \mathcal{G}^*$ and let $\deg_i^m$ denote the degree of i in G^m. At every round m, neighbouring agents always exchange monitoring information, and they follow punishment individual actions with a probability proportional to past deviations. Monitoring information includes reports and numbers of pending punishments. Specifically, for each round $m' < m$ and pair of agents (j, l), a report relative to m' and (j, l) specifies whether j interacted with l in round m', and whether l defected j. For each $c \in \{1 \dots n\}$ and agent j, agents keep the number of pending punishments to be applied to j in periodic rounds $(c + kn)_{k \geq 0}$. Before interacting with j in round m, i determines whether j should be punished. For this, i updates the number x of pending punishments for the period that includes m. If $m \leq n$, then $x = 0$. Otherwise, let x' be the previous number of pending punishments resulting from an identical update prior to round $m-n$. Given the round-$m-n$ reports, i determines $\deg_j^{m-n}$. Then, i sets x to $\max(0, x' - \deg_j^{m-n})$, and adds $\deg_j^{m-n}$ iff j defected some neighbour in $m-n$. i punishes j with probability $\min(1, x/\deg_j^m)$. (This is where the knowledge of degree comes into play.) After the interaction, i emits a report indicating that the interaction occurred in round m and signalling whether j defected i. i also updates its monitoring information basing on the information sent by its neighbours. For each report relative to round $m' < m$ not older than $m - n + 1$ and pair (k, l), if i does not have a report relative to m' and (k, l) and receives a new report from $j \neq l$, then i stores this report. In addition, for each period c not including m, i updates the number of pending punishments relative to every $l \neq i$ and c to the maximum between its value and the value sent by every neighbour $j \neq l$, capping it to lie in $\{0 \dots n-1\}$.

We present in Alg. 1 the pseudo-code for the strategy σ_i^{gen} of agent i. We use two variables: AP and RP. For each round m, i keeps a number AP_j^m of pending punishments to be applied to j in rounds $m, m+n, m+2n, \dots$, initially equal to 0 and never larger than $n-1$. Also, i keeps a report $\text{RP}_{jl}^m \in \{\mathit{Good}, \mathit{Bad}, \perp\}$ per pair of nodes (j, l) and round r signalling whether in round m: (i) j did not interact with l ($\perp$); (ii) l interacted and defected j (*Bad*); or (iii) l interacted and did not defect j (*Good*). RP_{jl}^m is initialized to $\perp$. Notice that these variables can be implemented by a finite state machine: at each round m, agents only need to store and forward information relative to each tuple (i, j, c) for $c \in \{0 \dots n-1\}$, corresponding to RP_{ij}^{m-c}, and information relative to each pair (i, c) for $c \in \{0 \dots n-1\}$, corresponding to AP_i^{m-c}. For the sake of exposition, we opt to not represent the state in this compact form.

This definition has the following properties when agents follow $\vec{\sigma}^{\text{gen}}$ at every round-m history h: (1) i cannot influence monitoring information that determines punishments applied to i, (2) we match each defection of i in round m to at least one punishment in future rounds, (3) a defection in round m triggers additional punishments to be applied in rounds $m+n, m+2n \dots$, and (4) the delay of additional punishments is bounded by $O(n^2)$. (1) ensures that i does not gain from lying about monitoring information. (2) ensures that, even if i saves the cost 1 of sending messages, i loses at least $\beta > 1$. (3) and (4) guarantee that this loss is discounted to the present by a lower bounded factor δ^{n^2}. This implies that, if agents are sufficiently patient (i.e., δ is sufficiently close to 1), then i prefers not to defect.

Theorem 7 shows that $\vec{\sigma}^{\text{gen}}$ enforces accountability with connectivity and sufficiently patient agents.

Theorem 7. *If the adversary is restricted by connectivity and agents are sufficiently patient, then $\vec{\sigma}^{gen}$ enforces accountability in general pairwise exchanges.*

Proof. We show that $\vec{\sigma}^{\text{gen}}$ is a $\mathcal{G}^*$-OAPE in general pairwise exchanges. Fix $G \in \mathcal{G}^*$, agent i, round-

Algorithm 1 σ_i^{gen}

for all j, m **do**
  $\text{AP}_j^m \leftarrow 0$
  for all l **do**
    $\text{RP}_{jl}^m \leftarrow \perp$

Upon round m
  for all neighbour j **do**
    $pr \leftarrow \min(1, \text{AP}_j^m / \deg_j^m)$
    With probability pr, cooperate ▷ Punishment
    Otherwise, punish ▷ Cooperation
End

After round m
  for all neighbour j **do**
    if j defects i in m **then**
      $\text{RP}_{ij}^m \leftarrow \textit{Bad}$
    else
      $\text{RP}_{ij}^m \leftarrow \textit{Good}$
      for all $m' \in \{m-n+1 \dots m-1\}$ **do**
        $\text{AP}_j^{m'} \leftarrow$ maximum between $\text{AP}_j^{m'}$ and values sent by each $l \neq j$
        Cap $\text{AP}_j^{m'}$ to lie in $\{0 \dots n-1\}$
        for all $l \neq i$ **do**
          if $\text{RP}_{lj}^{m'} = \perp$ and some $k \neq j$ sent $\text{RP}_{lj}^{m'}|_k \neq \perp$ **then**
            $\text{RP}_{lj}^{m'} \leftarrow v \neq \perp$ deterministically selected among received values
  for all agent $j \neq i$ **do**
    if $m \geq n$ **then**
      $\deg \leftarrow \#\{l \neq j \mid \text{RP}_{lj}^{m-n+1} \neq \perp\}$ ▷ Degree $\deg_j^{m-n+1}$
      $\text{AP}_j^{m+1} \leftarrow \max(0, \text{AP}_j^{m-n+1} - \deg)$
      if exists $l \neq j$ such that $\text{RP}_{lj}^{m-n+1} = \textit{Bad}$ **then**
        $\text{AP}_j^{m+1} \leftarrow \text{AP}_j^{m+1} + \deg$
End

m, round-m information set $I_i \in \mathcal{I}_i(G)$, $h \in I_i$, and actions a_i^*, a_i' such that $\sigma_i^{\text{gen}}(a_i^* | I_i) > 0$. Let $\Delta = \mathbb{E}^{\vec{\sigma}^1}(u_i \mid h, G) - \mathbb{E}^{\vec{\sigma}^2}(u_i \mid h, G)$, where $\vec{\sigma}^1 = \vec{\sigma}^{\text{gen}}|_{I_i, a_i^*}$, $\vec{\sigma}^2 = \vec{\sigma}^{\text{gen}}|_{I_i, a_i'}$, and $\vec{\sigma}^{\text{gen}}|_{I_i, a_i}$ represents the protocol that differs from $\vec{\sigma}^{\text{gen}}$ exactly in that i deterministically follows a_i at I_i. By the One-Shot-Deviation principle, it suffices to show that $\Delta \geq 0$ to prove the result. For every $m' \geq m$, let $\Delta^{m'}$ be the difference between the expected utility of i in round m', between agents following $\vec{\sigma}^1$ and $\vec{\sigma}^2$, i.e.,

$$\Delta^{m'} = \sum_r Pr^{\vec{\sigma}^1}(r \mid G, h) u_i(\vec{a}^{m'}) - \sum_r Pr^{\vec{\sigma}^1}(r \mid G, h) u_i(\vec{a}^{m'}),$$

where $\vec{a}^{m'}$ is the round-m' action profile in r. We prove six facts first, which hold for every run of $\vec{\sigma}^1$ and $\vec{\sigma}^2$ in G, conditioned on h. We say that a fact holds after i follows $\vec{a}_i'$ and $\vec{a}_i^*$ if the fact holds for all runs r^1 and r^2 of $\vec{\sigma}^1$ and $\vec{\sigma}^2$, respectively, such that $Pr^{\vec{\sigma}^1}(r^1 \mid G, h) > 0$ and $Pr^{\vec{\sigma}^2}(r^2 \mid G, h) > 0$. Fix agents j and l different from i, and rounds m', m''. Let $\text{RP}_{ij}^m|_{l,m'}$ and $\text{AP}_i^m|_{l,m'}$ be the values of RP_{ij}^m and $\text{AP}_i^m|_{l,m'}$, respectively, deterministically held by l at the end of round m' in every run of the considered strategies:

(a) Fact 1: If i follows a_i', then $\text{RP}_{ij}^m|_{l,m'}$ is accurate: (1) if i and j do not interact or it does not hold that $(j, m) \rightsquigarrow_i^G (l, m'+1)$, then the report says that they did not interact; (2) if i and j interact and $(j, m) \rightsquigarrow_i^G (l, m'+1)$, the the report says that i defected j iff i defected j in a_i'.

Proof. At the end of round m, the report of l says that i and j interacted iff $l = j$ and i and

j are neighbours. This is exactly the case when $(j,m) \rightsquigarrow_i^G (l, m+1)$. Moreover, (2) follows directly by construction of $\vec{\sigma}^{\text{gen}}$. In every subsequent round m', agent l updates $\text{RP}_{ij}^m|_{l,m'}$ to $\text{RP}_{ij}^m|_{l,m'+1}$ only if $\text{RP}_{ij}^m|_{l,m'} = \perp$ and some round-m' neighbour o sent $\text{RP}_{ij}^m|_{o,m'} \neq \perp$. By the hypothesis, $\text{RP}_{ij}^m|_{o,m'}$ is accurate, so it follows that $\text{RP}_{ij}^m|_{l,m'+1}$ is also accurate. □

(b) Fact 2: If $m' = m + n - 1$ and i follows a_i', then $\text{AP}_i^{m'+1}|_{l,m'} = y + \max(x - \deg_i^m, 0)$, where x is the maximum of $\text{AP}_i^m|_{o,m}$ for all agents $o \neq i$, and $y = \deg_i^m$ if i defects a neighbour in a_i' or $y = 0$ otherwise.

Proof. It is easy to show using induction that for every round $m'' \in \{m \dots m'\}$, $\text{AP}_i^m|_{l,m'} = x^{m''}$, where $x^{m''}$ is the maximum of $\text{AP}_i^m|_{o,m}$ for all agents $o \neq i$ such that $(o,m) \rightsquigarrow_i^G (l, m''+1)$. This is clearly true for $m'' = m$, since only l is causally influenced by l between m and $m+1$. For every other m'', l updates $\text{AP}_i^m|_{l,m''}$ to the maximum of the values sent by round-$m''-1$ neighbours (Lines 1.13-22), which by the induction hypothesis is $x^{m''}$. Since the adversary is restricted by connectivity, l is causally influenced by every agent of i without interference from i between m and $m+n$, so $\text{AP}_i^m|_{l,m'} = x$. After updating $\text{AP}_i^m|_{l,m'}$, l computes the value deg, which by Fact 1 is $\deg_i^m$, sets $\text{AP}_i^{m'+1}|_{l,m'} = \min(0, x - \deg_i^m)$, and adds $\deg_i^m$ if some report indicates that i defected some neighbour, which by Fact 1 is true iff i defected some neighbour in a_i'. This proves Fact 2. □

(c) Fact 3: If $m'' \neq m$ and $m' \geq m$, then the values of $\text{AP}_i^{m''}|_{l,m'}$ and $\text{RP}_{ij}^{m''}|_{l,m'}$ when i follows a_i' are the same as when i follows a_i^*.

Proof. If $m' = m$ and $m'' < m$, then those values depend only on the values received by l from agents different from i, which are determined by the history h. Continuing inductively, for every $m' > m.$, l updates $\text{AP}_i^{m''}|_{l,m'}$ and $\text{RP}_{ij}^m|_{l,m'}$ according to the values sent by round-m' neighbours $o \neq i$ of l, which are the same by the hypothesis, so the resulting values of $\text{AP}_i^{m''}|_{l,m'}$ and $\text{RP}_{ij}^m|_{l,m'}$ are the same, whether i follows a_i' or a_i^*. If $m'' > m$, then l does not update $\text{AP}_i^{m''}|_{l,m'}$ and $\text{RP}_{ij}^m|_{l,m'}$ when $m' < m''$. When $m' = m'' - 1$, l updates $\text{AP}_i^{m''}|_{l,m'}$ basing on the values $\text{AP}_i^{m''-n}|_{l,m'}$ and $\text{RP}_{ij}^{m''-n}|_{l,m'}$, which are the same by the proof of the case $m'' < m$. When $m' = m''$, l updates $\text{RP}_{ij}^{m''-n}|_{l,m'}$ iff l interacts with j in round m'', and updates it to *Good*, whether i follows a_i' and a_i^*, since i follows σ_i^{gen} after m and does not defect l. Finally, when $m' > m''$, the same arguments as above prove the result. □

(d) Fact 4: If $m' > m$, then the value $\text{AP}_i^{m''}|_{l,m'}$ is at least as high when i follows a_i' as when i follows a_i^*.

Proof. First, consider that $m'' < m$. By Fact 3, the values are the same whether i follows a_i' or a_i^*. Consequently, the same holds for all rounds $m'' + cn$ with $c > 0$. Regarding $m'' = m$, by Fact 2, the value $\text{AP}_i^{m''}|_{l,m'}$ is higher by $\deg_i^m$ when i follows a_i' than a_i^* if i defects some neighbour in a_i', or is the same otherwise. Again, the result follows directly from Fact 3 for all rounds $m + cn$ with $c > 0$. □

(e) Fact 5: If $m' > m$, then $\Delta^{m'} \geq 0$.

Proof. In round m', both i and its neighbours either cooperate or punish each other, so the benefits and communication costs are the same, whether i follows a'_i or a^*_i. By Fact 4, i is never punished with higher probability when i follows a^*_i than when i follows a'_i, hence the expected costs of punishments must be at least as high when i follows a'_i as when it follows a^*_i. This implies that $\Delta^{m'} \geq 0$. □

(f) Fact 6: If i defects some neighbour in a'_i, then $\sum_{m<m^*<m+n^2} \Delta^{m^*} \geq \beta \deg_i^m$.

Proof. Let $m^c = m + cn$ for $c \geq 0$, and let $\mathrm{AP}_i^{m^c}|_{o,m^c-1,a^*_i}$ and $\mathrm{AP}_i^{m^c}|_{o,m^c-1,a'_i}$ be the value of $\mathrm{AP}_i^{m^c}|_{o,m^c-1}$ when i follows a^*_i and a'_i, respectively. By Fact 2, if i defects some neighbour in a'_i, we have $\mathrm{AP}_i^{m^1}|_{o,m^1-1,a'_i} - \mathrm{AP}_i^{m^c}|_{o,m^1-1,a^*_i} > \deg_i^m$. For every $c \geq 0$ with $\mathrm{AP}_i^{m^c}|_{o,m^c-1,a^*_i} \geq \deg_i^{m^c}$, we have

$$\mathrm{AP}_i^{m^{c+1}}|_{o,m^{c+1}-1,a^*_i} - \mathrm{AP}_i^{m^c}|_{o,m^c-1,a^*_i} = \deg_i^{m^c}.$$

Let $m^* = m^c$ be the first round where $\mathrm{AP}_i^{m^c}|_{o,m^c-1,a^*_i} \leq \deg_i^{m^c}$. When i follows a^*_i, in round m^*, i is punished by an expected number of $\mathrm{AP}_i^{m^*}|_{o,m^*-1,a^*_i}$ neighbours, while not being punished afterwards in rounds $m^* + cn$ for $c > 0$. When i follows a'_i, i is punished in rounds $m^* + cn$ for $c \geq 0$ by an additional expected number of $\deg_i^m$ neighbours. Moreover, since i always has at least one neighbour and i never defects after m, by Fact 2 and the fact that $\mathrm{AP}_i^{m+n}|_{o,m+n-1,a'_i} \leq n-1$, we have $\mathrm{AP}_i^{m^n}|_{o,m^n-1,a'_i} = 0$. This implies that the additional punishments when i follows a'_i must occur in rounds between $m+1$ and $m+n^2$. By the same arguments as in the proof of Fact 5, it follows that the expected utility of i decreases at least by $\beta \deg_i^m$ in these rounds when i follows a'_i instead of a^*_i, as we intended to prove. □

We can now show that $\Delta \geq 0$. If i does not defect a neighbour in a'_i, then $\Delta^m \geq 0$ and by Fact 5 $\Delta \geq 0$. If i defects neighbours in a'_i, then $\Delta^m \geq -\deg_i^m$, and by Fact 6, we have $\Delta \geq -\deg_i^m + \delta^{n^2} \beta \deg_i^m$. So, if $\beta > 1$ and δ is sufficiently close to 1, then $\Delta \geq 0$. This concludes the proof. □

6.5 Avoiding Knowledge of Degree

We now discuss one type of interactions in general pairwise exchanges and corresponding implementation of $\vec{\sigma}^{\mathrm{gen}}$, assuming only that agents know the identities of their neighbours. We need at least two communication phases, so that neighbours i and j may first reveal their degree in phase 1, and then exchange values in phase 2 and punish each other accordingly. Unfortunately, two phases is not enough because i may lie about the degree. In particular, in phase 1, i may declare a higher degree than the real one to decrease the probability of being punished by j. We can address this by including in the report information about the declared degrees, and then punish agents that lie about their degrees. This is still not sufficient though, because of the following scenario. Suppose that i only has one neighbour j and only one pending punishment. If i does not lie, then j punishes i with probability 1 due to having a single pending punishment. If i declares a degree of $n-1$ instead, then the probability of being punished by j is only $1/(n-1)$. This yields an increase in the expected benefits from 0 to $\beta(n-2)/(n-1)$. Suppose that i also defects j by omitting messages in phase 2. The expected future loss must be at most β, to avoid the problem identified in the proof of need for eventual distinguishability. Since we only assume that $\beta > 1$, the loss may be lower than the gain. The problem is that i sends the degree before incurring the cost of sending messages in phase 2. We can avoid this by using a technique of delaying gratification employed in [18]. We need three communication phases. In phase 1, agents exchange monitoring information and the values ciphered with random private keys. In phase 2, they reveal the degrees. Finally, in phase

3, they decide whether to cooperate by sending the private keys, or to punish by sending arbitrary keys. Let α^κ be the cost of sending a key, such that the cost of sending phase 1 and 2 messages plus α^κ is 1. Agent i is punished for defecting j in phases 1 or 2 by not receiving the value in phase 3. In this case, i avoids at most $1+\alpha$, but loses $\beta > 1+\alpha$ (notice that, by consistency of beliefs, i always believes that it will receive the value in phase 3 prior to this phase, even if j does not send its value correctly ciphered, because i cannot distinguish ciphered values from arbitrary ones). The maximum utility gain obtained by i is when i has only one neighbour j, i lies to j by saying that its degree is $n-1$, and then defect j in phase 3. The maximum gain is $\beta(n-2)/(n-1)+\alpha^\kappa$, whereas the future loss is $\delta\beta$. If $\alpha^\kappa < \beta/(n-1)$, then the gain is less than β, and the future loss outweighs the gain for a δ sufficiently close to 1. Therefore, we have a protocol that enforces accountability.

6.6 Complexity

The bit complexity of $\vec{\sigma}^{\text{gen}}$ is $O(n^3)$: each message carries n^3 reports, with one report per pair of agents and round in $1 \cdots \rho$, carries n^2 accusations, and carries n^2 numbers of pending punishments. The maximum delay of punishments is $O(n^2)$. We can improve the complexity by assuming further restrictions on $\mathcal{G}^*$ and on the computational ability of agents. Specifically, the factors n^3 and n^2 are a function of (1) the maximum delay n of disseminating information relative to an agent, and (2) the maximum number n of agents relative to which an agent has to forward information. (1) can be improved by considering more restrictive assumptions about $\mathcal{G}^*$. As discussed in the definition of adversary restricted by connectivity, we only need that the neighbours of any given agent i in round m can causally influence all the neighbours of i in round $m+\rho$ for some constant ρ. If the evolving graphs satisfy locality properties that ensure that the neighbours of agent i's neighbours are likely to be i's neighbours in the near future, then ρ can be significantly smaller than n. Then, the bit complexity becomes $O(\rho n^2)$ and the maximum delay of punishments becomes $O(\rho n)$. Small-world networks such as social networks, which have a high clustering coefficient, are well known examples that satisfy such locality properties [12, 25].

Interestingly, locality properties also allow us the improve on (2). Agents only have to forward information relative to an agent i and round m in the ρ rounds following m. This is necessary to ensure that critical information reaches the agents capable of carrying a punishment. If during every period of ρ rounds each agent were only causally influenced by a limited number c of agents, then agents only needed to forward information relative to c agents, and the bit complexity and delay of punishments would be $O(\rho c^2)$ and $O(\rho c)$, respectively. Unfortunately, such reduction on the amount of information each agent forwards introduces a congestion problem. Specifically, we cannot let agents sending messages of varying size, since that would incentivize them to always forward the minimum required information to save bandwidth. Therefore, they must forward messages of fixed size, which implies that agents may have to discard information if they cannot fit all received information in a single message. This gives the opportunity for an agent i to deliberately generate false reports (or accusations) that flood the network, causing other agents to discard accusations against i. We can address this issue by appending to each report and accusation a signature of the issuer. A report of an interaction between i and j has issuer i, so a valid report of this interaction must contain a signature of i. The same applies to accusations and numbers of pending punishments. The only disadvantage of this approach is that it requires agents to be computationally bounded.

7 Discussion

Our results provide technical insight on the type of dynamic networks that can support pairwise exchanges in equilibria strategies. The need for timely punishments means that accountability cannot be enforced in some dynamic networks such as certain overlays for file-sharing. If connectivity is ensured, accountability may be enforced, even for general exchanges. Although for this

scenario we have assumed knowledge of degree prior to interactions, this can be relaxed if agents can exchange multiple messages per round. Finally, if exchanges are valuable, timely punishments are enough to enforce accountability, which opens the door to protocols that enforce accountability in a wide variety of dynamic networks. There are multiple open questions to be addressed in future work. It would be interesting to prove a stronger condition that would close the gap between an adversary restricted by distinguishability and one restricted by connectivity. Another open issue is collusion. We believe that both the necessary and sufficient conditions presented in this paper could be strengthened by generalizing the notion of causal influence without interference from individual agents to the absence of interference from members of a coalition. Given these conditions, the $\mathcal{G}^*$-OAPE strategies are resilient to collusion.

References

[1] I. Abraham, D. Dolev, R. Gonen, and J. Halpern. Distributed computing meets game theory: robust mechanisms for rational secret sharing and multiparty computation. PODC'06, pages 53–62. ACM, 2006.

[2] I. Abraham, D. Dolev, and J. Halpern. Distributed protocols for leader election: A game-theoretic perspective. DISC'13, pages 61–75. Springer-Verlag, 2013.

[3] Y. Afek, Y. Ginzberg, S. Landau Feibish, and M. Sulamy. Distributed computing building blocks for rational agents. PODC'14, pages 406–415. ACM, 2014.

[4] B. Cohen. Incentives build robustness in bittorrent. P2PEcon'03, Berkeley, CA, USA, June 2003.

[5] A. Fabrikant, A. Luthra, E. Maneva, C. H. Papadimitriou, and S. Shenker. On a network creation game. PODC'03, pages 347–351. ACM, 2003.

[6] I. Fainmesser. Community structure and market outcomes: A repeated games in networks approach. *American Economic Journal: Microeconomics*, 4(1):32–69, 2012.

[7] I. Fainmesser and D. Goldberg. Cooperation in partly observable networked markets repeated games played in a network. UFAE and IAE working papers, Brown University, 2012.

[8] M. Feldman, K. Lai, I. Stoica, and J. Chuang. Robust incentive techniques for peer-to-peer networks. EC'04, pages 102–111, New York, NY, USA, 2004. ACM.

[9] R. Guerraoui, K. Huguenin, A. Kermarrec, M. Monod, and S. Prusty. Lifting: lightweight freerider-tracking in gossip. Middleware'10, pages 313–333. ACM, 2010.

[10] J. Halpern and V. Teague. Rational secret sharing and multiparty computation: Extended abstract. STOC'04, pages 623–632. ACM, 2004.

[11] E. Hendon, H. Jacobsen, and B. Sloth. The one-shot-deviation principle for sequential rationality. *Games and Economic Behavior*, 12(2):274–282, 1996.

[12] P. W. Holland and S. Leinhardt. Transitivity in structural models of small groups. *Small Group Research*, 2(2):107–124, 1971.

[13] S. Jun and M. Ahamad. Incentives in bittorrent induce free riding. P2PECON'05, pages 116–121. ACM, 2005.

[14] D. Kreps and R. Wilson. Sequential equilibria. *Econometrica*, 50(4):863–894, 1982.

[15] F. Kuhn, N. Lynch, and R. Oshman. Distributed computation in dynamic networks. STOC'10, pages 513–522. ACM, 2010.

[16] D. Levin, K. LaCurts, N. Spring, and B. Bhattacharjee. Bittorrent is an auction: analyzing and improving bittorrent's incentives. In *Proceedings of the ACM SIGCOMM 2008 conference on Applications, technologies, architectures, and protocols for computer communication*, SIGCOMM'08. ACM, August 2008.

[17] H. C. Li, A. Clement, M. Marchetti, M. Kapritsos, L. Robison, L. Alvisi, and M. Dahlin. Flightpath: obedience vs. choice in cooperative services. OSDI'08, pages 355–368. Usenix Association, 2008.

[18] H. C. Li, A. Clement, E. L. Wong, J. Napper, I. Roy, L. Alvisi, and M. Dahlin. BAR gossip. OSDI'06, pages 191–204. Usenix Association, 2006.

[19] G. Mailath and L. Samuelson. *Repeated Games and Reputations*. Oxford University Press, 2007.

[20] T. Moscibroda, S. Schmid, and R. Wattenhofer. On the topologies formed by selfish peers. PODC'06, pages 133–142. ACM, 2006.

[21] G. Neglia, G. Presti, H. Zhang, and D. Towsley. A network formation game approach to study bittorrent tit-for-tat. NET-COOP'07, pages 13–22. Springer-Verlag, 2007.

[22] M. Osborne and A. Rubinstein. *A course in game theory*. The MIT Press, 1994.

[23] X. Vilaca and L. Rodrigues. Accountability in dynamic networks. Arxiv preprint arXiv:1602.03885, May 2016.

[24] X. Vilaça and L. Rodrigues. On the range of equilibria utilities of a repeated epidemic dissemination game with a mediator. ICDCN '15, pages 19:1–19:10. ACM, 2015.

[25] D. J. Watts and S. H. Strogatz. Collective dynamics of small-world networks. *Nature*, 393(6684):440–442, 1998.

[26] E. Wong and L. Alvisi. What's a little collusion between friends? PODC'13, pages 240–249. ACM, 2013.